Are Vaccines Fungible? Regression Discontinuity Evidence from a Large Aid Program

Sarah Dykstra, Amanda Glassman, Charles Kenny, & Justin Sandefur*

This version: October 17, 2014†

Abstract

Since 2001, a single international aid consortium has accounted for over half of total vaccination expenditure in over 70 developing countries, delivering roughly $7 billion worth of vaccines and complementary funding to countries with GNI below a strict threshold, originally set at $1,000 per capita. We exploit this cutoff in a regression discontinuity (RD) framework. In middle-income countries near the threshold, results suggest large inflows of free vaccines led to a modest, statistically insignificant increase in vaccination rates across a range of diseases. There is little evidence of waste. Instead, it appears a majority of free vaccines were delivered to children who would have been vaccinated anyway, displacing domestic immunization efforts. We emphasize that our RD estimates do not rule out positive impacts for the poorest countries farther from the eligibility threshold.

Keywords: aid, vaccination, immunization, fungibility, regression discontinuity

JEL Classification Codes: F35 - Foreign Aid; H51 - Government Expenditures and Health; I15 - Health and Economic Development; O11 - Macroeconomic Analyses of Economic Development

*All authors are affiliated with the Center for Global Development (CGD) in Washington D.C. We are grateful for helpful comments from Victoria Fan, Willa Friedman, and numerous seminar participants. CGD does not take institutional positions, and the views expressed here do not necessarily reflect those of CGD's board or its funders.

†The most recent version of the paper can be downloaded at www.cgdev.org/vaccinesrd. A web appendix containing additional information on data construction and sample countries can be accessed at www.cgdev.org/vaccinesrdappendix.

1 Introduction

Vaccines are an obvious target for international aid given their proven efficacy in preventing disease and relatively low cost per dose. As of the mid 2000s, an estimated 634,000 children in Africa and Southeast Asia died annually of pneumococcal disease, 433,000 of rotavirus, and 278,000 of Haemophilus influenza type B (Watt et al., 2009; WHO/UNICEF, 2012a; O'Brien et al., 2009). In addition, over one million children in the 2000 birth cohort are predicted to die of hepatitis B (Goldstein et al., 2005). Effective vaccines already exist for all of these diseases, capable of averting anywhere from 74% of rotavirus deaths (Munos et al., 2010) up to 95% of hepatitis B deaths (Goldstein et al., 2005), with costs ranging from around $6 for the pneumococcal vaccine to as low as 34 cents per child for hepatitis B.

Beginning in 2000, a wide array of public and private aid donors pooled funds to increase immunization rates in poor countries through the primary vehicle of bulk purchase of vaccines. This partnership, known as the Global Alliance for Vaccines and Immunizations (GAVI), has cumulatively disbursed $7 billion worth of vaccines and financing to 77 countries with a combined population of 4.5 billion. GAVI aims to deliver sufficient new and underutilized vaccines in-kind each year to provide a full course of shots to every infant in a participating country.

The Alliance employs an income-based eligibility criteria for countries to receive free vaccines, which initially excluded all countries with a per capita gross national income (GNI) in excess of $1,000 in 1998 at market exchange rates. In this paper we exploit the unique eligibility rules for GAVI aid to estimate the impact of free vaccines on vaccination rates in countries near that income cutoff, i.e., whether vaccines delivered have translated into more children vaccinated. Our identification strategy is based on a regression discontinuity at the cutoff line. We show that this threshold has been fairly strictly enforced, providing a clean natural experiment in cross-country time-series data.

Results suggest that for countries near the threshold, free vaccines led to small but statistically insignificant increases in vaccination coverage between 2000 and 2013. This is true across all six vaccines for which data is available: hepatitis B, Haemophilus influenza type B (Hib), and diphtheria, pertussis, and tetanus (DPT) which are often combined in a single pentavalent vaccine funded by GAVI, as well as measles which received limited very

GAVI support, and two newer vaccines which GAVI has prioritized but which have achieved limited roll-out to date, pneumococcal disease and rotavirus.[1]

To explain the lack of large overall impacts near the cutoff, we quantify the role of waste and fungibility. We find little evidence of outright waste, defined as an excess of GAVI vaccines delivered to a country relative to the total number of infants vaccinated.[2] There is some evidence that waste is a transitory phenomenon: GAVI delivers vaccines sufficient to cover 100% of the population, and countries appear to require several years to ramp up vaccination rates to this level. But for the more established vaccines, vaccination rates approaching 90% imply limited waste.

Fungibility is a much larger factor than waste in explaining the limited impact of GAVI on vaccination rates for countries near the eligibility threshold. For Hib and DPT, where we cannot reject the null of zero impact, we can reject the alternative null of zero fungibility, i.e., aid did not increase vaccination rates by a number equivalent to the doses purchased by GAVI. Given the Alliance usually provides close to 100% of the vaccines required to immunize all children and that coverage rates for some of the vaccines GAVI supports were already at very high levels in many eligible countries in 2000, this result should be expected. Indeed, the GAVI model forced fungibility on countries. If they had not reduced their own purchases from 2000 levels, they would have had more doses than children to vaccinate.

Perhaps the most important caveat on these results is that we are reliant on fairly noisy, cross-national data. While the point estimates in our preferred specifications are close to zero, standard errors are large, and confidence intervals often span non-negligible effects from a welfare standpoint. Nevertheless, it is important to note that even with noisy data, the cleaner identification of causal effects provided by the regression discontinuity design allows

[1]Consistent with these vaccination results, we find no impact on infant or child mortality, although our power calculations indicate that finding such an effect would be unlikely even in a best-case scenario. As of 2000, about 44% of the roughly 9 million child deaths in developing countries were attributed to vaccine-preventable causes (Black et al., 2003). But diseases covered by GAVI vaccines constitute only 15% of those deaths (Watt et al., 2009; WHO/UNICEF, 2012a; O'Brien et al., 2009; Goldstein et al., 2005; Morris et al., 2008). Furthermore, given the limited roll-out of the newly-developed pneumococcal and rotavirus vaccines up to 2013, it is likely that any impact on child mortality observed during the 2000s would operate through the Hib vaccine, representing no more 3% of total child deaths in 2000 (Watt et al., 2009). Note that null results for child mortality neglect GAVI's potential impact on adult morbidity and mortality. However, GAVI's own estimates predict reduced adult mortality primarily via increased hepatitis B vaccination rates, a channel where we find no impact.

[2]This is an underestimate of waste if there is any domestic expenditure on the same vaccines. It is an overestimate of waste if vaccines are delivered to children outside the recommended age range covered in vaccination statistics, though as we discuss below, DHS survey evidence suggests the latter phenomenon is relatively rare.

us to reject the hypothesis that, absent GAVI, progress on immunization roll-out would have stalled in countries that were not eligible near the cut-off.

A second caveat regards impact in countries further from the GAVI cutoff line. There has been considerable convergence in vaccine coverage rates between countries above and below the cutoff since 2000, and the average rate of vaccination across all countries has been rising. Hepatitis B, Hib, and DPT vaccination rates on average approach 90% in countries with an active GAVI program by the end of our sample period in 2013. GAVI in-kind support may well have played a large role in raising vaccination rates in the poorest countries above what they would have been absent the program.

A final caveat relates to the impact that GAVI may have had on prices of new vaccines in the market in general, that may have allowed countries near the cutoff to access lower prices for new vaccines than would have been the case in GAVIs absence. Historical price data for vaccines for every country in our sample is not available in the public domain.

Our contribution to the literature is threefold. First, we provide an independent measure of the causal impact of GAVI support. Previous evaluations of this $7 billion program have lacked credible counterfactuals. GAVI itself estimates it has immunized 440 million additional children and averted 6 million future deaths (GAVI Alliance, 2014b). This calculation assumes that in the absence of GAVI aid, vaccination rates would have remained constant since 2000 – an assumption our RD allows us to test and reject for countries just outside the eligibility threshold. A more sophisticated set of independent, academic analyses relies on dynamic panel-data GMM estimators proposed by Blundell and Bond (1998) to estimate GAVI's impact. Using this approach, Lu et al. (2006) find evidence of a positive, significant impact of GAVI funding on DPT vaccination rates for countries with initial coverage rates under 65%. In contrast, Hulls et al. (2010) extend this model to include additional years of data and find no such effect below 65% baseline DPT coverage, but a significant impact for countries between 65% and 80%. These findings are subject to the standard concerns about casual inference using off-the-shelf GMM panel-data estimators (Roodman, 2009). In addition, the justification for partitioning the sample using baseline DPT coverage is not clear. The partitioning highlights and may exacerbate the challenge of causal inference: because almost all countries with very low baseline DPT rates were within GAVI's income eligibility threshold, these models are estimating GAVI's impact by comparing aid recipients (compliers) to countries who failed to negotiate a GAVI aid package for a variety of

reasons (non-compliers). Almost all variation in treatment status is driven by unobserved characteristics of recipients.[3]

Our RD estimates provide a more credible measure of GAVI's impacts, relying on exogenous features of the aid program for identification. An obvious caveat is that the local average treatment effect we estimate is specific to the immediate vicinity of the GAVI eligibility cut-off, i.e., for the wealthiest GAVI recipients. There is *a priori* justification to believe the impact of free vaccines will be smallest for these recipients, both because baseline immunization rates are higher, and because these middle-income countries have the greatest capacity to purchase new vaccines as they emerge. However, focusing on countries near the eligibility threshold remains of considerable policy relevance as GAVI seeks a new round of financing from major donors, where negotiations have focused on revising the eligibility rules, particularly 'graduation' rules for rapidly growing economies.

Second, a recurrent debate around aid impact involves fungibility. Donors may be attracted to aid for vaccines not only by their apparent low cost and high efficacy, but by the perception that in-kind distribution mitigates fungibility concerns. Previous work has found significant fungibility of aid generically (Devarajan et al., 1999; Feyzioglu et al., 1998; Pack and Pack, 1993),[4] although with little attention to the endogeneity of aid. Recent findings of high levels of fungibility for health aid specifically (Lu et al., 2010) have been criticized both on identification grounds (Roodman, 2012) and misspecification due to the conflation of on-budget and off-budget aid (de Sijpe, 2013). Our regression discontinuity framework provides a cleaner empirical test of fungibility, and focuses squarely on the type of in-kind, off-budget aid designed to overcome it.

Third, we address a gap in the large cross-country literature on aid effectiveness, which has been overwhelmingly concerned with aid's effects on economic growth, producing mixed and mostly modest results (Burnside and Dollar, 2000; Easterly et al., 2004; Rajan and

[3]Note that Lu et al. (2006) and Hulls et al. (2010) also focus on just one of several vaccines funded by GAVI. We expand our focus beyond DPT to include hepatitis B, pneumococcal disease, rotavirus, and Hib, as well as impacts on infant and child mortality. The pneumococcal and rotavirus vaccines were only widely introduced in the late 2000s, and we take advantage of more recent data to estimate GAVI's impact on their diffusion.

[4]Pack and Pack (1990) is an exception to the general pattern of results, finding no evidence of fungibility in sector-specific aid to Indonesia. Van de Walle and Mu (2007) is an exception to the otherwise limited attention to causal identification in this literature, finding fungibility within (but not across) broad sector categories in the context of a field experiment.

Subramanian, 2008; Clemens et al., 2012).[5] This growth focus fits awkwardly with the stated goals of health aid, whose growth effects may be limited (Acemoglu and Johnson, 2007; Clemens et al., 2012). Many prominent commentators in the aid debate, including both advocates and skeptics of international aid as a whole, appear optimistic about the effectiveness of health aid (Deaton, 2013; Easterly, 2014). But while the empirical support for this optimism is evident in impact evaluation of specific aid-funded programs, the relationship between large aid programs operating in multiple countries and health impact is not well-established in the literature. A limited number of studies find salutary effects of health aid on child mortality (Arndt et al., 2014; Mishra and Newhouse, 2009), but again causal inference in these dynamic panel regressions is open to debate (Roodman, 2009). This paper seeks to provide additional evidence on the link between aid and health outputs and outcomes.

The remainder of the paper is organized as follows. Section 2 describes the structure of the GAVI aid program and documents our data sources. Section 3 describes our econometric framework to measure standard treatment effects as well as fungibility and waste. Section 4 describes the results. Section 5 concludes.

2 Program background and data sources

2.1 Eligibility

In GAVI's first board meeting, attendees agreed on a threshold for support. "Countries with small resources and a lack of purchasing power have been considered to be in greatest need of financial support for the new vaccines. It is proposed that this be initially interpreted as those with a GNP/capita equal to or less than 1,000 USD" (GAVI Alliance, 1999). No detailed justification was offered for this threshold, other than that it was "IDA-like", i.e., similar to the cut-off for concessional World Bank loans explored by Galiani et al. (2014), an issue we return to below. Since then, GAVI's own financial sustainability models provide some ex post rationalization for the cutoff based on countries' ability to pay for vaccines.

[5] It is interesting to note that one of the most recent notable contributions to the aid and growth literature is the most methodologically similar to our analysis of vaccines, and finds quite large, positive, growth effects from aid. Galiani et al. (2014) construct a regression discontinuity analysis by exploiting the abrupt income eligibility threshold for International Development Assistance (IDA), the World Bank's most concessionary loan window.

This eligibility rule implies that the appropriate forcing variable in our analysis is the World Bank's measure of gross national income (GNI) per capita (Atlas method, current US$). In phase 1, the $1,000 GNI per capita threshold was measured using 1998 data. Seventy-five countries fell below this cut-off and were eligible to apply for support throughout the first phase.[6] In phase 2, the threshold was maintained at $1,000 based on 2003 GNI data and the list of eligible countries was updated (GAVI Alliance, 2009).[7] Beginning in 2011, the eligibility threshold was increased to $1,500 and adjusted annually for inflation in subsequent years (GAVI Alliance, 2014a). The cut-off was increased to $1,500 because it is roughly equivalent to the original threshold in 2000, adjusting for inflation, and because in a constrained economic climate, it made sense for GAVI to retain fewer eligible countries (GAVI Alliance, 2009). The number of eligible countries dropped from 72 in phase 2 to 55 in the last two phases.

Because some GNI estimates found in the World Bank's databank have been revised in the time since they were used to determine GAVI eligibility, we draw GNI data from original versions of the World Bank's World Development Indicator publications. Section A.6 in the web appendix reports GNI data and vaccination rates for the countries in our sample by phase. Summary statistics appear in Table 2.

As we show below, GAVI remained fairly faithful to its initial, somewhat arbitrary income-eligibility rule. Unfortunately, from a research perspective, it has largely ignored any changes in country eligibility due to changes in the income threshold or economic growth. Once a country has qualified for GAVI aid, subsequent loss of eligibility is unlikely to put an end to aid flows, as we show in Section A.2 and Table 2 in the appendix.

In addition to the income threshold, GAVI instituted coverage thresholds which determined eligibility for participation in the new and under-used vaccines support (NVS) and immunization services support (ISS) program. In phase 1, countries with a DPT vaccination rate of less than 80% were eligible to apply to receive ISS support, which provided

[6] Our sample for phase 1 contains 71 eligible countries. Bolivia had a GNI of $1,010 in 1998 and was not eligible until 2002, when its GNI dropped to $990 per capita. Timor-Leste did not become a country until 2002, at which point it became eligible. No GNI data were available Cuba and Djibouti in 1998, lower-middle income countries.

[7] GAVI's program dates can be confusing. We define phases in terms of eligibility criteria, which differs slightly from how the phases are defined on the website. GAVI's phase 1 operations extended from 2000 until 2005 and phase 2 programming began in 2006, the year in which eligibility was updated based on 2003 GNI data (see Azerbaijan and the Democratic Republic of Congo's applications for funding in 2005). GAVI's website lists phase 2 as beginning in 2007 because the strategic goals governing phase 2 were not approved until that year. Phase 3 began in 2011 for both programming and strategic goals, but because the eligibility threshold was updated annually at this point, we define 2012 as phase 4.

performance-based financing for increased DPT coverage. This threshold was eliminated in phase 2, allowing all GAVI-eligible countries to apply. Countries with DPT coverage rates greater than 50% were able to apply for NVS funding, support grants for vaccine purchases. This threshold was increased to 70% in phase 3. In practice, these non-income eligibility thresholds have very little predictive power as to which countries received GAVI assistance, as discussed in Section A.2 in the appendix.

2.2 Aid flows

Most GAVI funding goes to purchasing vaccines, about $4.5 billion of its total $6.1 billion in disbursements from 2000 to 2013. These vaccines can be roughly divided into two groups: older, cheaper vaccines which dominated GAVI's programming for the first several years of its existence, and newer, more expensive vaccines which have rolled out more recently. The first group includes the vaccines for diphtheria, pertussis, and tetanus (DPT), Hib, and hep B, which GAVI has more recently replaced with a single pentavalent vaccine that covers all five diseases. The cost per dose for DPT and hep B vaccinations is well under $1, as shown in Table 1, while the pentavalent vaccines costs between $1.30 and $3.50 per dose. Spending on these vaccines accounts for roughly 45% of GAVI's total budget. GAVI also spent $1.6 billion on broader health systems support, auto-disable syringes and immunization systems support (ISS), a performance-based program designed to incentivize increased DPT coverage.

The second group includes the vaccines for rotavirus and pneumococcal disease, both of which were essentially non-existent in the developing world prior to GAVI roll-out in the late 2000s. These vaccines are more costly, at $3.50 to $5.00 per dose for rotavirus and $3.50 to $7.00 for the pneumococcal vaccine. These vaccines have absorbed roughly a fifth of GAVI's total budget, or $1.2 billion, concentrated in later years.

In recent years, GAVI constitutes the majority of foreign aid for vaccinations and spends more per capita on vaccines in eligible countries than countries spend themselves – by a 6:1 ratio in 2012. The WHO and UNICEF Joint Reporting Form has collected country reports of government and non-government spending on immunizations since 2006. On average from 2006 to 2010 (GAVI phase 2), countries that were eligible for GAVI aid spent about $0.16 per capita per year on immunization programs, and received about $0.36 (mostly in kind) per capita from other sources. By 2012 (GAVI phase 4) government spending had fallen to $0.09 while non-government spending had risen to about $0.45. (See Table 2 and Figure 2.)

These numbers are roughly consistent with GAVI's own records, which show that eligible countries received, on average, $0.26 per capita from GAVI in 2006-10 and $0.60 in 2012.

GAVI does not report on the number of doses purchased by country from 2000 to 2012. We construct an estimate of the number of doses of each vaccine purchased using data from GAVI and UNICEF, which we use for the waste calculations in subsequent parts of the paper. UNICEF reports on the total number of doses of each vaccine purchased worldwide and the total amount spent on each vaccine, which we use to calculate the average cost per dose across all formulations of a vaccine.[8] While GAVI does not provide information vaccine presentations purchased, they do publish data on spending by vaccine. We use the average cost per dose to estimate how many doses of each vaccine GAVI purchased by country and year. When using this data to estimate vaccination rates had GAVI not purchased vaccines, we account for buffer stocks and wastage using vaccine-specific rates reported by countries in the Joint Reporting Form, using the mean wastage rate by vaccine when an estimate is not available.[9]

2.3 Vaccination outcomes

In the years since the launch of GAVI, there has been considerable progress in vaccination rates in countries that were below the income threshold. Figure 1 displays vaccination rates according to WHO data for GAVI-eligible countries at phase 1 and those developing countries above the $1,000 threshold using a fixed sample of 71 eligible and 73 ineligible countries. As can be seen, for DPT, hepatitis B and Hib there has been considerable convergence in coverage rates between the two groups, driven by very rapid growth of coverage in GAVI-eligible countries. The coverage gap between the two sets of countries has fallen from 20 percentage points to 10 points for DPT, 38 to 11 points for hepatitis B and 15 to 7 in the case of Hib all as vaccination rates have increased in both sets of countries. The progress in vaccine rollout in eligible countries may or may not reflect the impact of GAVI, of course. Progress before GAVI's launch, and progress in non-eligible countries both suggest other

[8] These estimates were compared to UNICEF's awarded prices per dose for all presentations of a vaccine, and our results are within the range of costs per dose paid by UNICEF, available here.

[9] Wastage rates are self-reported through the Joint Reporting Form. As with country estimates of immunization coverage, it is possible that these numbers are biased or inaccurate (Ronveaux et al., 2005). Bosch-Capblanch et al. (2009) find that a majority of countries were able to correctly calculate wastage in their second data quality audit, a standard WHO process used to verify vaccination data quality.

factors have had a role to play. This paper will examine the relative importance of GAVI support in raising immunization rates for countries near the income threshold.

The immunization coverage indicator we use is the WHO/UNICEF Estimates of National Immunization Coverage (WUENIC) for 2000 to 2013. These coverage estimates are available for all countries on an annual basis and measure the proportion of infants surviving to their first birthday who have completed 6 vaccine series: DPT, hepatitis B, Hib, measles, pneumococcal and rotavirus. We also construct an average index of DPT, hepatitis B, and Hib ("penta"), which are commonly administered together in a single pentavalent vaccine.

WHO constructs estimates using two sources of coverage data, administrative and survey. Survey data is considered the most reliable measure of coverage, but is only available intermittently and for a limited set of countries (Lim et al., 2008). Administrative data are drawn from reports of vaccinations performed by service providers, which are then aggregated to the national level and reported annually to the WHO through the Joint Reporting Form. Administrative coverage figures are subject to both numerator and denominator bias due to erroneous reporting of doses delivered to children older than 12 months, under-reporting of doses provided by private clinics and inaccurate target population size, especially in cases where population projections are based on outdated censuses (Burton et al., 2009).

To account for these issues, WHO and UNICEF review coverage estimates for each country and vaccine on an annual basis to identify and address situations where data may be compromised and are not accurate representations of true coverage (WHO/UNICEF, 2012b). Estimates from any recent household surveys (primarily DHS and MICS) are compared to country estimates.[10] WHO consults with local officials, considers potential biases in each source and attempts to construct a pattern over time which reconciles different estimates.

These estimates are based in part on administrative data, which have the potential to over-report coverage levels due to measurement error and incentives from performance-based payment systems (Sandefur and Glassman, 2014; Murray et al., 2003; Lessler et al., 2011). Lim et al. (2008) observe large discrepancies between survey estimates and official country estimates of coverage, particularly among countries receiving ISS funding, where over-reporting increased significantly after the introduction of the program. They find that official estimates

[10]There is often a delay of several years before results from surveys like the DHS are available. If new data from surveys become available which improve upon previous years' estimates, earlier estimates are also corrected. Contextual information, such as stock-outs or other disruptions to service delivery, is also considered. (Mitchell et al., 2012)

of DPT coverage show more rapid improvement than survey-based estimates. WHO estimates are a measurable improvement on country estimates and exhibit less over-reporting, but they are still higher than survey estimates on average (Lim et al., 2008).

Because participation in ISS is associated with inflated coverage rates, a potential consequence of using WHO estimates is that we detect an effect on DPT vaccination rates where no true effect exists. While over-reporting due to ISS payments is a concern, the results of Lim et al. (2008) suggest that any misreporting would bias our results in favor of finding a positive effect of GAVI funding on immunization rates. The incentive to misreport created by the ISS did not extend to vaccines other than DPT.

Lastly, it is important to note that WHO vaccination rates refer to the target population, which for all of the vaccines examined here is defined as infants under the age of one year. A potential threat to our research design would arise if vaccines were routinely administered to children above this age threshold, a phenomenon referred to as "catch-up vaccination." This would be especially problematic if free vaccines delivered by GAVI led to increased catch-up vaccination in the treatment group, leading to an underestimate of aid's impact. Fortunately for the research design, there is strong evidence this is not the case.

While we cannot observe catch-up vaccination in our data, the phenomenon is well-documented in the literature. Clark and Sanderson (2009) draw on Demographic and Health Survey (DHS) data from 45 countries to measure the delays in completing DPT vaccination courses. Delays of several weeks are quite common in many countries, but these delays very rarely push the date of completion beyond the one-year mark. In the full sample of 45 countries, 50% of children had received the final dose of DPT by 20 weeks of age, and 95% by 30 weeks.

3 Empirical strategy

3.1 Potential outcomes framework

Suppose GAVI gives Ghana enough doses of pentavalent vaccine to immunize 91% of Ghanaian infants, as was the case in 2013 (Government of Ghana, 2013).[11] Where do these vaccines

[11] GAVI provided 100% of vaccines in 2013 but required the Government of Ghana to provide $0.20 of a total cost of $2.18 per dose in cofinancing (Government of Ghana, 2013)

end up? First, some vaccines will be used to vaccinate children who would not otherwise have been vaccinated (which is what we mean when we refer to impacts). Second, some vaccines may be used to vaccinate children who would have been vaccinated even in the absence of GAVI aid (our measure of fungibility). Third, some vaccine doses will simply not be used (which we term waste). Our measures of impact, fungibility, and waste add up to the total number of vaccines delivered.

To construct these measures, we require a clear counterfactual scenario. To explain how the RD design provides the necessary counterfactual to measure both impacts and fungibility, we first present our empirical strategy in the standard potential outcomes framework, following Lee and Lemieux (2009).

Suppose that for each observation i the outcome takes a value of $Y_i(1)$ if i is eligible for treatment and $Y_i(0)$ if it is not. Country i is eligible for aid if its income X_i is below the threshold c. For each observation we observe only one of these potential outcomes, and we can estimate $E[Y_i(0)|X > c]$ for the ineligible countries above the GAVI cutoff, and $E[Y_i(1)|X < c]$ for the eligible countries below the cutoff.

The key assumption underlying RD estimation is that potential outcomes are continuous in X, the forcing variable (Hahn et al., 2001). Under this assumption, the treatment effect at the cutoff becomes:

$$E[Y_i(1) - Y_i(0)|X = c] = \lim_{\substack{\varepsilon \to 0 \\ <}} E[Y_i(1)|X = c + \varepsilon] - \lim_{\substack{\varepsilon \to 0 \\ >}} E[Y_i(0)|X = c + \varepsilon] \qquad (1)$$

This expression makes clear that RD estimates yield a *local* average treatment effect (LATE) in a very specific sense: the LATE applies to observations in the immediate vicinity of the cutoff, c, as ε approaches zero. This LATE is an intent-to-treat (ITT) effect, as we have defined treatment purely in terms of treatment eligibility. Below we consider the effect of treatment on the treated (TOT), i.e., the LATE applicable to countries near the threshold who actually received GAVI aid.

Now turn to the question of measuring fungibility and waste. Let D_i denote the total vaccine doses delivered by GAVI to country i, measured as a proportion of the target-age population the aid could cover in an ideal scenario. GAVI often calculated its aid so that $D_i = 100\%$, especially for hepatitis B, Hib, and DPT. As our starting point, we assume

that if vaccines are completely non-fungible and there is zero waste, then aid should raise vaccination rates by an amount equivalent to the total vaccines received. This implies that

$$E[\tilde{Y}_i(1) - \tilde{Y}_i(0)|X = c] = 0, \tag{2}$$

where $\tilde{Y}_i(1) \equiv Y_i(1) - D_i(1)$. In words, once we subtract GAVI vaccine doses from the vaccination rate, we should observe no difference between eligible and ineligible countries if aid is not fungible and is not wasted.

To make this concrete, return to our numerical example. Ghana ($1420 per capita GNI in 2011) and Nicaragua ($1,600) fell just on either side of the eligibility threshold for GAVI's phase 5. Ghana recorded a hepatitis B vaccination rate in 2013 of 90%, compared to 98% for Nicaragua, while Ghana received free vaccines equivalent to 91% coverage and Nicaragua received none. Equation (2) implies that if there is no fungibility or waste, we should expect Ghana's vaccination rate to be equal to Nicaragua's after subtracting the aid to Ghana – 90% coverage minus 91% aid (i.e., -1%) should equal 98%. The gap between -1% and 98% implies that GAVI aid was either wasted or crowded out domestic expenditure in Ghana.

3.2 Regression specification

Our basic empirical strategy is a straightforward application of regression discontinuity methods: we regress changes in vaccination and mortality rates within a country on a function of the baseline per capita GNI variable used to determine GAVI eligibility, and test for a discontinuous jump in outcomes at the threshold. We test for balance by showing that pre-treatment trends in vaccination exhibited no discontinuity at the threshold. We estimate our core regression both parametrically (using a linear regression with polynomial terms on either side of the cut-off) and non-parametrically, reporting both intent-to-treat (ITT) and average treatment on the treated (TOT) effects. For the latter, we use a fuzzy regression discontinuity design to account for the fact that many GAVI-eligible countries did not receive aid in a given year – though almost all eligible countries received aid in at least some years.

They key challenge in estimation is to provide a reasonable empirical approximation of the (continuous) function linking the forcing variable to the outcome. We follow the advice of Gelman and Imbens (2014) and focus on low-order polynomials on either side of the cut-off, varying the bandwidth of data used in estimation.

To measure the intent-to-treat effect, we regress change in vaccination and mortality rates (Y) in country c and year t on a p^{th}-order polynomial of log GNI per capita in year t, X_{ct}, an eligibility dummy indicating that the country fell below the GAVI eligibility threshold in the given year, E_{ct}, and the interaction of the treatment dummy and the polynomial terms.

$$\begin{aligned} Y_{ct} &= \beta_{00} + \beta_{01}\tilde{X}_{ct} + \beta_{02}\tilde{X}_{ct}^2 + \ldots + \beta_{0p}\tilde{X}_{ct}^p \\ &\quad + \rho E_{ct} + \beta_{11}E_{ct}\tilde{X}_{ct} + \beta_{12}E_{ct}\tilde{X}_{ct}^2 + \ldots + \beta_{1p}E_{ct}\tilde{X}_{ct}^p + \eta_{ct} \end{aligned} \tag{3}$$

In practice, we present linear and quadratic results, $\{p : 1, 2\}$. Note that the tilde over the GNI terms indicates that they have been de-meaned, such that $\tilde{X}_{ct} \equiv X_{ct} - \bar{X}_{ct}$, enabling us to interpret the ρ coefficient as an estimate of the treatment effect at the discontinuity.

Because estimates may be sensitive to the bandwidth chosen, we employ the method proposed by Imbens and Kalyanaraman (2012) to select the optimal bandwidth to minimize mean-squared error. The optimal bandwidth selection is implemented using the "rd" package in Stata, documented in Nichols (2007) and subsequent updates. Rather than allowing the bandwidth to vary as we estimate effects on, say, hepatitis B and DPT, we use a benchmark bandwidth that is the rough midpoint of the Imbens and Kalyanaraman (2012) estimates across the six vaccines studied. We also report results using half and twice this bandwidth.

In addition to the ITT effect, we are also interested in the effect treatment on the treated (TOT). This is a specific LATE estimate, relevant to eligible countries which may have endogenously selected into treatment and, in this instance, only such countries which are near the eligibility threshold. We can estimate the TOT effect in a two-stage least squares framework commonly referred to as a fuzzy regression discontinuity design.

In the second stage, we regress Y_{ct} on a p^{th}-order polynomial of log GNI per capita in year t, X_{ct}, the predicted value of the treatment dummy indicating that the country received GAVI funding in the given year, T_{ct}, and the interaction of the predicted treatment dummy and the polynomial terms.

$$\begin{aligned} Y_{ct} &= \beta_{00} + \beta_{01}\tilde{X}_{ct} + \beta_{02}\tilde{X}_{ct}^2 + \ldots + \beta_{0p}\tilde{X}_{ct}^p \\ &\quad + \rho\widehat{T}_{ct} + \beta_{11}\widehat{T_{ct}\tilde{X}_{ct}} + \beta_{12}\widehat{T_{ct}\tilde{X}_{ct}^2} + \ldots + \beta_{1p}\widehat{T_{ct}\tilde{X}_{ct}^p} + \eta_{ct} \end{aligned} \tag{4}$$

In the first stage, we regress the actual GAVI treatment dummy on an indicator of GAVI eligibility, which takes a value of one if the country's GNI was above the cut-off in that

phase, E_{ct}. In addition, we control for the same polynomial function of log GNI as used in the second stage in equation 4, as well as the interaction of these polynomial terms and the eligibility dummy.

$$\begin{aligned} T_{ct} &= \gamma_{00} + \gamma_{01}\tilde{Y}_{ct} + \gamma_{02}\tilde{Y}_{ct}^2 + \ldots + \gamma_{0p}\tilde{Y}_{ct}^p \\ &\quad + \pi E_{ct} + \gamma_{11}E_{ct}\tilde{Y}_{ct} + \gamma_{12}E_{ct}\tilde{Y}_{ct}^2 + \ldots + \gamma_{1p}E_{ct}\tilde{Y}_{ct}^p + \varepsilon_{ct} \end{aligned} \tag{5}$$

Note that we rely on a binary indicator as the treatment variable, T, although dollar figures and dosage counts are both available from GAVI, because no instrument is available to explain variation in the size of GAVI aid packages across eligible countries.

4 Results

4.1 First stage results and balance tests

Before turning to the main results, we address two important preliminaries: (a) ensuring that GAVI aid exhibits a discontinuous jump at the eligibility threshold, and (b) demonstrating that other variables which were pre-determined at the time of GAVI's launch do not.

The first-stage regressions show that GAVI's sharp eligibility threshold has been enforced quite strictly. Results of estimating equation (5) are reported in Table 3 and Figure 3. The probability of receiving GAVI aid falls by about 7% for each additional log point of per capita GNI, and jumps more than 50% at the threshold, with minor variation across specifications. This jump is significant at the 1% level with linear and quadratic controls for log per capita GNI, and is fully robust to allowing the GNI function to vary on either side of the threshold. The jump is significantly less than 100% primarily because of non-compliance rather than defiance: after GAVI was launched in 2000, disbursements began only gradually in many countries, and many country-year cells are coded as zero, particularly in earlier years.

In theory, GAVI eligibility depended not only on the income threshold, but also on separate thresholds defined by baseline DPT vaccination rates, as described in Section 2. In practice, it appears these thresholds were not strictly enforced. A country's classification based on baseline DPT shows some relationship with the probability of receiving GAVI aid, but this relationship is not robust and only marginally significant in the best cases. We explore this issue in detail in the appendix (see Table 3). Because baseline DPT adds

little or no predictive power to the first-stage regression, we focus exclusively on the income threshold.

The sharp jump in GAVI aid at the GNI threshold raises the question of whether countries differed previously at this point, and whether GAVI simply chose a pre-existing discontinuity in country performance. For the variables where we have data on pre-treatment trends – including vaccination rates for hepatitis B, Hib, DPT, and measles, as well under-5 and under-1 mortality rates – we fail to reject the null of parallel pre-treatment trends.

More formally, regression discontinuity designs are subject to balance tests, analogous to those reported for randomized experiments. The assumption that counterfactuals are a continuous function of the forcing variable implies, in our setup, that baseline values of the outcome variables (circa 2000) should exhibit no jump at the GAVI eligibility threshold. To test this, we apply equation 3 to baseline values, and report results in Table 4. As seen in the table, sample sizes are relatively small (roughly 10% the size of our main specifications for the treatment effects) and thus power is somewhat limited. Nevertheless, we fail to detect any significant differences in trends in vaccination or mortality in 2000.

4.2 Treatment effects on vaccination and mortality rates

Our treatment effect estimates are uniformly underwhelming. We find no significant impact of free vaccines delivered by aid donors on vaccination rates across any of the six diseases tested.

Intent-to-treat (ITT) results from the reduced form regression shown in equation (3) are reported in Table 5. Each column reports results for a different outcome indicator, and the three panels present results using various bandwidths (i.e., ranges of data as defined by the forcing variable). We treat the linear specification with a bandwidth of 1 log point per capita GNI as our preferred specification. Results from this specification are largely consistent with the wider bandwidth and quadratic specifications.

In the benchmark specification, point estimates are round 1% with standard errors between 2% and 6%. The one exception among the six diseases is Hib, with a point estimate of 10%, though this is also insignificant and not robust across specifications. Across the various diseases, specifications and bandwidths, the results are fairly evenly spread around zero. The quadratic specification with the narrowest bandwidth (0.5 log GNI) produces some

anomalous results – in particular, a point estimate implying GAVI reduced Hib vaccination rates 28% – but the standard errors here are very large, possibly indicating we're trying to estimate too many parameters with too few data points when fitting a quadratic to a thin slice of data on either side of the threshold.

To appreciate the magnitude of these point estimates, it is useful to examine the treatment-on-the-treated (TOT). Two-stage least squares results based on the fuzzy regression discontinuity framework of equation (4) are reported in Table 7. As we found no significant ITT effects in Table 5, by construction there are no significant TOT effects to report in Table 7. The point estimates from the TOT estimates highlight, though, that the confidence interval on GAVI's impact in countries where it is operational spans some reasonably large magnitudes from a policy perspective. The point estimate for hep B is just 1.8% in the linear IV specification (bw=1), but the 95% confidence interval includes positive effects upwards of 20%.

4.3 Fungibility

Finding no evidence of an impact on vaccination rates is, of course, not the same as finding evidence of no impact. One might worry that our non-results reflect a lack of power to detect important effect sizes. To address this concern, we explore whether our data, which fails to reject the null of zero impact, can reject two relevant alternative hypotheses.

First, we show that vaccination rates around the GAVI threshold are not consistent with the hypothesis that free vaccines are not fungible (and not wasted). Second, and closely related, we show that the data is also inconsistent with the idea that all 440 million additional children vaccinated with GAVI support would have remained unvaccinated absent the program. Consider these in turn.

As explained in Section 3, the regression discontinuity framework provides a way to measure the fungibility of aid. Under the null hypothesis of no fungibility and no waste, vaccination rates in eligible and non-eligible countries on either side of the eligibility threshold should be indistinguishable after subtracting the free vaccines from the former, as shown in equation (2).

In our preferred specification (linear, bandwidth of 1 log point) in Table 8, there is significant evidence of fungibility or waste for both hepatitis B and DPT. This result is

particularly strong and robust for hepatitis B, where vaccination rates in eligible countries are 56% lower than they should be in the absence of fungibility or waste. For DPT, we find that vaccination rates are about 30% below where they would be without fungibility or waste, though this result disappears with the narrower 0.5 log point bandwidth.

Notably, we find no evidence of fungibility or waste for four of the six vaccines, including rotavirus and pneumococcal disease which were introduced later and more slowly, and measles (MCV) which has received limited GAVI support. Given the uniformly null impacts on vaccination rates across all diseases documented in the previous subsection, the difference in the fungibility and waste results across diseases here is largely driven by differences in the quantity of aid delivered. No impact with no aid is no evidence of fungibility or waste. In contrast, finding no impact given the very large deliveries of aid for hepatitis B and DPT constitutes strong evidence of fungibility or waste.

Second, setting aside the issue of fungibility and waste, we turn to a different counterfactual. We can reject the hypothesis that all 440 million additional children vaccinated with the support of GAVI since 2000 would have been left unvaccinated absent GAVI in-kind support that vaccination rates would have remained at baseline levels.[12] Recall that the core assumption of regression discontinuity analysis stated in Section 3 is that potential outcomes are a continuous function of the forcing variable. This implies that baseline vaccination rates in GAVI eligible countries (circa 2000) should be equal to current vaccination rates in ineligible countries near the threshold – i.e., we should observe zero increase in vaccination in our control group.

This suggests the following regression. Adding time subscripts to our notation, we define the counterfactual in the data, Y_i^s as follows:

$$Y_i(0) = Y_{it}^s = \begin{cases} Y_{i,baseline} & \text{if } X_{it} \leq c; \\ Y_{it} & \text{if } X_{it} > c. \end{cases}$$

The key point is that both Y^a and Y^s should be continuous functions of the forcing variable, X, with no jump at the GAVI threshold. If Y^a or Y^s is significantly lower on the eligible side of the cutoff, this implies that GAVI's official impact claims are biased upward by the use of an incorrectly low counterfactual. Results are reported in Table 4 in the appendix

[12]This is defined as the absolute number of children vaccinated the year before GAVI began funding a vaccine in a country.

4.4 Waste

Up to now we have lumped together fungibility and waste, but they have very different implications from a public finance perspective. Fungibility may be optimal from a recipient country's perspective – as free vaccines liberate domestic resources for other legitimate purposes – even if it frustrates the objectives of the aid donor. Waste, in which free vaccines simply go unused, is obviously suboptimal for all parties. Waste is a particular concern with in-kind aid programs such as GAVI. Money doesn't spoil, but vaccines do.

While we cannot observe waste directly, in some extreme cases, we can infer it indirectly when aid donors provided more free vaccines than the total number of children vaccinated in a country. Thus we propose a conservative measure of waste, defined as the excess of aid over and above the total vaccination rate in a given country year:

$$\text{Waste}_{it} \equiv \max[-(Y_{it} - D_{it}), 0]. \tag{6}$$

For example, suppose GAVI delivered sufficient doses of pentavalent vaccines to immunize 97% of infants in Djibouti against Hib in 2011, yet Djibouti recorded a Hib vaccination rate of just 87%. We interpret this as an indication that vaccines sufficient to cover 10% of Djiboutian infants were wasted. This estimate is conservative inasmuch as the government Djibouti (or private Djiboutian households) may have also bought some vaccines from domestic resources, pushing the implied wastage rate even higher.

There are three main assumptions required for our estimates of waste. The first is that vaccines are used only for the target population, i.e., infants under one year of age. As discussed in Section 2, there is strong evidence from the Demographic and Health Surveys that this is indeed the case (Clark and Sanderson, 2009). The second is that our constructed measure of GAVI-purchased doses accurately reflects the number of doses provided by GAVI. As prices per dose of vaccine varied slightly across vaccine presentation, waste may be over- or under-estimated depending on the exact price paid in each country.

We must also make some assumption about the time period over which vaccines apply. There is some evidence of lumpiness in GAVI deliveries – i.e., countries receiving in excess of 100% of their needs in one year, and significantly less in the following year. In this case, equation (7) will overestimate waste. To avoid any overestimation of waste, we again take a very conservative stance, and average all aid receipts and vaccination rates within countries

over all the years in our data set in which GAVI was active, 2001-2013. In effect, we replace equation (7) with

$$\text{Waste}_{it} \equiv \max[-(\bar{Y}_i - \bar{D}_i), 0]. \tag{7}$$

By averaging over years, this implicitly allows that vaccines never spoil and are saved for future years, and assumes only that they are not reallocated across countries once they are disbursed.

For each of six vaccines, a scatterplot is shown in Figure 5 with the average doses delivered per annum in terms of the population under 1 (D_i) for each country plotted again average vaccination rate.[13] The solid black line is a 45-degree line. Points below the line represent countries which received more doses from GAVI on average than the number of children under 1 vaccinated over the same years. For hepatitis, for instance, Ghana received enough doses per annum to vaccinate 83% of children, but had an average vaccination rate of 75% during the same time period. Across all countries receiving hepatitis B vaccines from GAVI, the number of doses lost to waste was sufficient to vaccinate approximately 19% of the population – though we would emphasize that in many cases waste occurred where vaccination rates were approaching 100%, so this reflects a global inefficiency, but often a simple arithmetic inevitability within countries. For other diseases, the doses delivered and waste rates are considerably lower, as seen in the figure.

So far we have focused on average waste level across all aid recipients. To compare these numbers to our estimates of impact and fungibility, we need to know the waste levels for a representative country at the GAVI eligibility threshold. To calculate this waste level, we repeat our benchmark RDD regression specification, using waste as the dependent variable. Results are reported in Table 9. In our preferred specification (linear, bandwidth of one log point of GNI), waste rates are statistically indistinguishable from zero for four of the five diseases calculated. (The same would be true for measles, which we do not report because GAVI disbursed so few measles doses that waste is not a significant issue.) The one exception is Hib, where we estimate that 6% of doses were lost to waste for a country near the threshold. Some care is required in interpreting this result, as in most cases, GAVI distributed hepatitis B, Hib, and DPT as a bundle through the pentavalent vaccine. Notably, when we average the vaccination rates across these three vaccines and compare the pentavalent deliveries, we find no evidence of waste at the threshold.

[13]The sample includes only GAVI eligible countries that received greater than zero doses of the vaccine in question in each year (i.e., compliers).In the figure, the percent of dose series purchased by GAVI is capped at 100%.

4.5 Summary decomposition

We are now prepared to complete the decomposition described in the introduction. The econometric estimates from the previous sections allow us to divide the free vaccines provided by GAVI's aid program between those that immunized children who wouldn't have otherwise have been vaccinated (impact), vaccines that immunized children who would have been vaccinated anyway (fungibility), and vaccines that were apparently never administered to anyone (waste). We estimate each piece in a separate regression, and the results presented in Table 10 are not strictly guaranteed to add up, but for the most part come fairly close.

To see how the decomposition works, imagine a hypothetical country whose income happens to fall just inside the GAVI eligibility threshold, and which is average in every respect given its income. This hypothetical threshold country would have received enough vaccines from GAVI each year from 2001 to 2013 to vaccinate 58% of its infants against hepatitis B, 30% against Hib, 36% against DPT, and just a few percent against rotavirus and pneumococcal disease. Where did these doses go?

In the case of hepatitis B, the results in Table 7 showed that aid raised vaccination rates by just 1.8%. The results in Table 8 suggest that fungibility and waste together accounted for enough vaccines to cover about 56% of infants. And the estimates in Table 9 suggest that little, if any, of that was pure waste. All of these numbers are summarized in Table 10 which implies that for hepatitis B almost all of the free vaccines provided by aid were indeed given to children in need of vaccination (a victory) but that almost all of these children would have been vaccinated without aid, given the counterfactual countries on the other side of the cutoff.

The results are qualitatively similar for DPT and the average index of hepatitis B, DPT, and Hib. In each case our hypothetical country would have received sufficient vaccines to cover a sizeable share of infants, and there is modest and statistically insignificant evidence that these free vaccines raised the overall vaccination rate. Instead, a majority of doses crowded out domestic vaccination efforts. There is virtually no evidence of outright waste.

For Hib, there is still no significant evidence of a positive impact, but the point estimates are much larger – representing a TOT effect of nearly 18 percentage points. The point estimate for fungibility is zero.[14] In essence, the data is too noisy to provide a firm indication

[14]The positive impact, zero fungibility, and non-zero waste estimates for Hib are internally inasmuch as the estimated Hib coverage rate for countries just outside the eligibility threshold are approximately zero. Any measured vaccination coverage for eligible countries near the threshold is counted as an impact. Any excess of aid over coverage is counted as waste rather than fungibility since the counterfactual is roughly zero coverage.

of where the free vaccines ended up, but point towards a non-trivial improvement in coverage. This is consistent with the fact that baseline coverage in 2000 for Hib was much lower than for hepatitis B and DPT. GAVI's decision to prioritize the pentavalent vaccine in lieu of all three of those vaccines might simultaneously have had no impact on hepatitis B or DPT coverage while pulling up Hib rates.

For the newer vaccines introduced in later GAVI phases, the decomposition is relatively uninformative. Little aid has been delivered on average over this period. The counterfactual for aid-recipients near the eligibility threshold is close to zero, but even so, that data cannot distinguish potential positive impacts of 1 to 3 percentage points from zero, or from outright waste.

Finally, as a validation of the data overall, the column corresponding to measles – which GAVI essentially did not fund, and we use as a placebo test – shows no impacts, no fungibility, and no waste.

5 Conclusion

To summarize the results, our main findings relate to the five diseases covered by the pentavalent vaccine (hepatitis B, Haemophilus influenza type B, diphtheria, pertussis, and tetanus), the single largest focus of GAVI's aid program. We find that the delivery of large quantities of free vaccines to developing countries with per capita incomes around $1,000 had no statistically significant impact on vaccination rates for these diseases. There is, however, statistically significant evidence of fungibility, in the sense that free vaccines appear to have displaced domestic vaccination spending in countries near the threshold. On the positive side, we find very little evidence of outright waste.

Our estimates based on a regression discontinuity design are broadly consistent with some previous panel data analyses (Lu et al., 2006), inasmuch as both find no significant impact of GAVI on vaccination rates for countries which are wealthier or had more developed vaccination systems at baseline. But our results contradict other previous panel data analyses (Lu et al., 2006) finding positive effects for countries with higher baseline vaccination rates.

It is important to emphasize, yet again, that all of our econometric estimates are applicable to countries near the aid eligibility threshold. There are *a priori* reasons to suspect the impact of free vaccines will be smallest near the threshold where our regression discontinuity

estimates apply. Free vaccines will be more likely to displace domestic vaccination programs in these middle-income countries, compared to low-income countries which may have weaker vaccination programs in the absence of aid. We also note that we have not estimated GAVI's impact on morbidity or other health outcomes via, for instance, the dissemination of safer vaccination techniques.

An additional caveat is that our estimates have focused on the impact of GAVI's direct aid to individual countries – its private transfers, as opposed to its public goods provision and activities that may have had substantial international spillovers. GAVI supported the development of a new vaccine, and it is not evident that the pneumococcal vaccine including serotypes prevalent in low-income countries would have existed by 2013 absent GAVI's advance market commitment. Additionally, GAVI's efforts may have had indirect impacts that extended beyond the strict eligibility threshold and are not captured here. Large aid purchases grew the market for vaccines and may have contributed to a drop in vaccine prices over time. The Alliance has a specific strategy to bring new entrants into the vaccine market and drive down prices. This may have had an effect on prices faced by non-GAVI countries. Between 2000 and 2010 the price of the hepatitis B vaccine dropped from US$ 0.56 to US$ 0.18 per dose, for example. GAVI also brought enormous international attention to the issue of vaccine coverage. This may have altered political calculations regarding vaccine spending and coverage above the threshold.

Nevertheless, our finding that free vaccines largely displaced funding for domestic vaccination programs in threshold countries may have broader relevance to the aid literature and policy discussions. This finding casts doubt on the aid strategy of using in-kind transfers to bypass concerns about fungibility, or of targeting only low-cost, easy-to-deliver interventions where the potential for duplication may be high. The design of GAVI's intervention may have essentially forced fungibility. If the Alliance provided enough pentavalent vaccines to fully immunize every child in the country, clearly countries would (and should) not purchase additional pentavalent vaccines even if they would have done so absent GAVI's intervention.

Turning to implications for GAVI's aid program in the future, our findings are open to a variety of interpretations. An obvious response would be to lower the income eligibility threshold for free vaccines that are cheap to buy, in order to avoid duplicating national efforts, since aid appears to have little or no marginal impact on vaccination rates near this threshold. Alternatively, GAVI might consider *raising* the GNI threshold to give even more countries access to the bulk purchasing discounts it receives through UNICEF, while phasing

out the actual aid transfer of free vaccines. The wisdom of these or other approaches hinges on feasibility considerations beyond the scope of this study.

References

Acemoglu, D. and S. Johnson (2007). Diseases and development: The effect of life expectancy on economic growth. Journal of Political Economy 115(6), 925.

Arndt, C., S. Jones, and F. Tarp (2014). Assessing foreign aid's long run contribution to growth and development. World Development.

Black, R. E., S. S. Morris, and J. Bryce (2003). Where and why are 10 million children dying every year? The Lancet 361(9376), 2226–2234.

Blundell, R. and S. Bond (1998). Initial conditions and moment restrictions in dynamic panel data models. Journal of econometrics 87(1), 115–143.

Bosch-Capblanch, X., O. Ronveaux, V. Doyle, V. Remedios, and A. Bchir (2009). Accuracy and quality of immunization information systems in forty-one low income countries. Tropical Medicine & International Health 14(1), 2–10.

Burnside, C. and D. Dollar (2000). Aid, policies, and growth. American Economic Review 90(4), 847–868.

Burton, A., R. Monasch, B. Lautenbach, M. Gacic-Dobo, M. Neill, R. Karimov, L. Wolfson, G. Jones, and M. Birmingham (2009). WHO and UNICEF estimates of national infant immunization coverage: methods and processes. Bulletin of the World Health Organization 87(7), 535–541.

Clark, A. and C. Sanderson (2009). Timing of children's vaccinations in 45 low-income and middle-income countries: an analysis of survey data. The Lancet 373(9674), 1543–1549.

Clemens, M. A., S. Radelet, R. R. Bhavnani, and S. Bazzi (2012). Counting chickens when they hatch: Timing and the effects of aid on growth. The Economic Journal 122(561), 590–617.

de Sijpe, N. V. (2013, December). The Fungibility of Health Aid Reconsidered. Journal of Development Studies 49(12), 1746–1754.

Deaton, A. (2013). The great escape: Health, wealth, and the origins of inequality. Princeton: Princeton University Press.

Devarajan, S., A. S. Rajkumar, and V. Swaroop (1999). What does aid to Africa finance? Mimeograph, Development Research Group, World Bank, Washington D.C.

Easterly, W. (2014). The Tyranny of Experts: Economists, Dictators, and the Forgotten Rights of the Poor. New York: Basic Books.

Easterly, W., R. Levine, and D. Roodman (2004). Aid, policies, and growth: Comment. American Economic Review, 774–780.

Feyzioglu, T., V. Swaroop, and M. Zhu (1998). A panel data analysis of the fungibility of foreign aid. The World Bank Economic Review 12(1), 29–58.

Galiani, S., S. Knack, L. C. Xu, and B. Zou (2014). The effect of aid on growth: Evidence from a quasi-experiment. World Bank Policy Research Working Paper 6865.

GAVI Alliance (1999). First board meeting minutes. http://libdoc.who.int/hq/1999/GAVI_99.02.pdf. Online; accessed 2014-08-06.

GAVI Alliance (2009). Review of the eligibility policy. http://www.gavialliance.org/About/Governance/GAVI-Board/Minutes/2011/16-Nov/Minutes/Review-of-the-Eligibility-Threshold/. Online; accessed 2014-06-21.

GAVI Alliance (2014a). Countries eligible for support. http://www.gavialliance.org/support/apply/countries-eligible-for-support/. Online; accessed 2014-06-21.

GAVI Alliance (2014b). Healthy children with a healthy future. http://www.gavialliance.org/Library/Publications/GAVI/GAVI-Alliance-replenishment-brochure/. Online; accessed 2014-08-06.

Gelman, A. and G. Imbens (2014). Why high-order polynomials should not be used in regression discontinuity designs. www.nber.org/papers/w20405. NBER Working Paper 20405.

Goldstein, S. T., F. Zhou, S. C. Hadler, B. P. Bell, E. E. Mast, and H. S. Margolis (2005). A mathematical model to estimate global hepatitis b disease burden and vaccination impact. International Journal of Epidemiology 34(6), 1329–1339.

Government of Ghana (2013). Gavi alliance annual progress report 2012. http://www.gavi.org/Country/Ghana/Documents/APRs/Annual-progress-report-Ghana-2012/. Online; accessed 2014-08-21.

Hahn, J., P. Todd, and W. Van der Klaauw (2001). Identification and estimation of treatment effects with a regression-discontinuity design. Econometrica 69(1), 201–209.

Hulls, D., P. Venkatachalam, K. Kumar, T. Cochrane, H. Kaur, N. Gulati, and D. Jones (2010). Second GAVI evaluation report. http://www.gavialliance.org/library/gavi-documents/evaluations/second-gavi-evaluation-2006-2010/. Online; accessed 2014-06-05.

Imbens, G. and K. Kalyanaraman (2012). Optimal bandwidth choice for the regression discontinuity estimator. The Review of Economic Studies 79(3), 933–959.

Lee, D. S. and T. Lemieux (2009). Regression discontinuity designs in economics. NBER Working Paper No. 14723.

Lessler, J., C. J. E. Metcalf, R. F. Grais, F. J. Luquero, D. A. Cummings, and B. T. Grenfell (2011). Measuring the performance of vaccination programs using cross-sectional surveys: a likelihood framework and retrospective analysis. PLoS medicine 8(10), e1001110.

Lim, S. S., D. B. Stein, A. Charrow, and C. J. Murray (2008). Tracking progress towards universal childhood immunisation and the impact of global initiatives: a systematic analysis of three-dose diphtheria, tetanus, and pertussis immunisation coverage. The Lancet 372(9655), 2031–2046.

Lu, C., C. M. Michaud, E. Gakidou, K. Khan, and C. J. Murray (2006). Effect of the Global Alliance for Vaccines and Immunisation on diphtheria, tetanus, and pertussis vaccine coverage: an independent assessment. The Lancet 368(9541), 1088–1095.

Lu, C., M. T. Schneider, P. Gubbins, K. Leach-Kemon, D. Jamison, and C. J. Murray (2010). Public financing of health in developing countries: a cross-national systematic analysis. The Lancet 375(9723), 1375–1387.

Mishra, P. and D. Newhouse (2009). Does health aid matter? Journal of Health Economics 28(4), 855–872.

Mitchell, V., V. J. Dietz, J. M. Okwo-Bele, and F. T. Cutts (2012). Immunization in developing countries. In S. Plotkin, W. Orenstein, and P. Offit (Eds.), Vaccines (Sixth edition ed.). Saunders.

Morris, S. K., W. J. Moss, and N. Halsey (2008). Haemophilus influenzae type b conjugate vaccine use and effectiveness. The Lancet infectious diseases 8(7), 435–443.

Munos, M. K., C. L. F. Walker, and R. E. Black (2010). The effect of rotavirus vaccine on diarrhoea mortality. International journal of epidemiology 39(suppl 1), i56–i62.

Murray, C. J., B. Shengelia, N. Gupta, S. Moussavi, A. Tandon, and M. Thieren (2003). Validity of reported vaccination coverage in 45 countries. The Lancet 362(9389), 1022–1027.

Nichols, A. (2007). Causal inference with observational data. Stata Journal 7(4), 507.

O'Brien, K. L., L. J. Wolfson, J. P. Watt, E. Henkle, M. Deloria-Knoll, N. McCall, E. Lee, K. Mulholland, O. S. Levine, and T. Cherian (2009). Burden of disease caused by streptococcus pneumoniae in children younger than 5 years: global estimates. The Lancet 374(9693), 893–902.

Pack, H. and J. R. Pack (1990). Is foreign aid fungible? The case of Indonesia. The Economic Journal, 188–194.

Pack, H. and J. R. Pack (1993). Foreign aid and the question of fungibility. The Review of Economics and Statistics, 258–265.

Rajan, R. G. and A. Subramanian (2008). Aid and growth: What does the cross-country evidence really show? The Review of Economics and Statistics 90(4), 643–665.

Ronveaux, O., D. Rickert, S. Hadler, H. Groom, J. Lloyd, A. Bchir, and M. Birmingham (2005). The immunization data quality audit: verifying the quality and consistency of immunization monitoring systems. Bulletin of the World Health Organization 83(7), 503–510.

Roodman, D. (2009). A note on the theme of too many instruments*. Oxford Bulletin of Economics and Statistics 71(1), 135–158.

Roodman, D. (2012). Doubts about the evidence that foreign aid for health is displaced into non-health uses. The Lancet 380(9846), 972–973.

Sandefur, J. and A. Glassman (2014). The political economy of bad data: Evidence from african survey and administrative statistics. CGD Working Paper 373.

Van de Walle, D. and R. Mu (2007). Fungibility and the flypaper effect of project aid: Micro-evidence for vietnam. Journal of Development Economics 84(2), 667–685.

Watt, J. P., L. J. Wolfson, K. L. O'Brien, E. Henkle, M. Deloria-Knoll, N. McCall, E. Lee, O. S. Levine, R. Hajjeh, K. Mulholland, et al. (2009). Burden of disease caused by haemophilus influenzae type b in children younger than 5 years: global estimates. The Lancet 374(9693), 903–911.

WHO/UNICEF (2012a). Estimated rotavirus deaths for children under 5 years of age. http://www.who.int/immunization/monitoring_surveillance/burden/estimates/rotavirus/en/. Online; accessed 2014-08-08.

WHO/UNICEF (2012b). User's reference to country reports of WHO and UNICEF estimates of national infant immunization coverage. http://www.childinfo.org/files/Immunization_WUENIC_guide_and_mark-up.pdf. Online; accessed 2014-06-21.

Table 1: GAVI vaccines, prices and purchases

Disease	Description	Presentations	Vaccine price	GAVI purchases[15]
Diphtheria, Pertussis, Tetanus (DPT)	Diphtheria is a potentially acute disease affecting respiratory tract, caused by a bacterial toxin, affecting mainly children. Pertussis is a potentially acute, highly contagious disease of the respiratory tract, caused by a bacteria, frequently complicates into pneumonia among children. Tetanus is an infectious bacterial disease causing muscular stiffness and spasms, primarily birth associated.	DPT, DTaP or pentavalent (TetraDTP-HepB-Hib)	$.09-.20 per dose for DPT presentation $1.30-3.50 per dose for pentavalent	DPT: N/A Incentivized through ISS support. Pentavalent: $2.35 billion
Haemophilus influenza type b (Hib)	Hib causes bacterial meningitis and an important cause of severe pneumonia in children <5 yrs old	Monovalent Hib, DPT-Hib, pentavalent vaccine (see DTP)	$1.58-3.20 per dose for DPT-Hib presentation	DPT-Hib: $11.2 million Pentavalent: $2.35 billion
Hepatitis B (Hep B)	Hepatitis B is a viral infection that attacks the liver and causes both acute and chronic disease. Most mortality occurs many years after the infection, therefore mainly in adulthood.	Monovalent hepatitis B vaccine, pentavalent (see DTP)	$.16-.80 per dose for monovalent presentation $.69-1.10 for DPT-HepB	DPT-HepB: $158.6 million Pentavalent: $2.35 billion
Rotavirus	Rotavirus causes severe, dehydrating diarrhea. Many strains of rotavirus exist.	RV1, RV5 (vaccines cover the referenced number of rotavirus strains)	$3.50-5.00 per dose	$187.6 million
Pneumococcal disease	S. pneumonia has >90 serotypes, and related infections cause meningitis, bacteraemia and pneumonia, and more common illnesses such as sinusitis and otitis media. More severe disease is called "invasive pneumococcal disease" (IPD).	PCV-7, PCV-10, PCV-13 (vaccines cover the referenced number of serotypes)	$3.50-7.00 per dose	$1.4 billion

[15]Cost per dose approximations are drawn from UNICEF awarded prices per dose, 2000-2012. GAVI purchase totals are from 2000-2013.

Table 2: Summary statistics

	Baseline		Phase 1		Phase 2		Phase 3		Phase 4		Phase 5	
	Eligible	Ineligible	Eligible	Ineligible	Eligible	Ineligible	Eligible	Ineligible	Eligible	Ineligible	Eligible	Ineligible
DPT (%)	69.5	89.3	74.1	89.6	80.0	92.3	80.9	90.7	81.6	91.3	80.6	90.8
	(21.0)	(12.1)	(19.0)	(12.3)	(15.9)	(9.28)	(16.6)	(9.25)	(15.9)	(9.33)	(16.0)	(11.1)
Hepatitis B (%)	16.8	54.9	41.9	70.6	72.2	89.0	77.6	90.1	79.3	90.7	79.7	90.3
	(32.5)	(43.8)	(40.2)	(37.0)	(29.5)	(18.6)	(23.2)	(11.6)	(21.3)	(11.4)	(19.2)	(11.9)
Hib (%)	3.87	18.6	11.3	41.1	40.5	64.0	73.6	74.9	75.0	79.6	77.1	83.6
	(18.3)	(34.3)	(28.5)	(44.1)	(41.9)	(41.9)	(30.0)	(33.0)	(27.9)	(30.3)	(23.4)	(24.8)
"Penta" (%)	30.1	54.3	42.5	67.1	64.3	81.8	77.4	85.3	78.7	87.2	79.1	88.2
	(18.6)	(21.0)	(23.4)	(23.1)	(24.0)	(17.1)	(21.7)	(14.5)	(20.2)	(13.6)	(18.3)	(13.1)
Rotavirus (%)	0	0	0	0	2.29	7.33	2.40	17.7	5.81	19.9	9.79	25.5
	(0)	(0)	(0)	(0)	(13.6)	(21.9)	(14.0)	(34.2)	(19.3)	(35.7)	(26.9)	(37.8)
Pneumococcal (%)	0	0	0	0	0.58	3.84	11.6	16.2	19.8	23.4	29.7	26.5
	(0)	(0)	(0)	(0)	(7.51)	(16.1)	(26.7)	(31.3)	(35.5)	(37.6)	(39.9)	(39.5)
Measles (%)	68.6	89.7	73.5	89.6	78.2	91.7	78.9	90.8	80	90.9	79.1	90.8
	(20.3)	(10.5)	(18.7)	(12.4)	(16.0)	(10.8)	(15.2)	(9.69)	(14.9)	(9.83)	(16.3)	(10.6)
Under-5 mortality rate	108.6	33.1	89.6	27.7	77.2	20.8	.	.	77.9	25.2	.	.
	(56.6)	(25.7)	(49.3)	(24.1)	(42.6)	(15.7)	(.)	(.)	(34.4)	(22.9)	(.)	(.)
Infant mortality rate	70.6	26.0	59.8	21.8	52.7	16.9	.	.	54.1	20.1	.	.
	(30.3)	(17.1)	(27.5)	(15.8)	(24.0)	(11.0)	(.)	(.)	(20.2)	(15.0)	(.)	(.)
GNI per capita ($)	430.8	2682.7	430.8	2682.7	479.4	2959.4	729.1	4652.7	778.1	4899.4	797.7	5053.0
	(212.0)	(1479.5)	(210.7)	(1470.6)	(230.5)	(1539.9)	(312.2)	(2317.7)	(328.7)	(2366.7)	(364.5)	(2399.3)
Gov't vaccine spending ($/capita)	.	.	.	.	0.17	1.14	0.093	1.14	0.096	1.09	.	.
	(.)	(.)	(.)	(.)	(0.23)	(1.33)	(0.097)	(1.19)	(0.094)	(1.03)	(.)	(.)
Total vaccine spending ($/capita)	.	.	.	.	0.42	1.46	0.56	1.55	0.53	1.53	.	.
	(.)	(.)	(.)	(.)	(0.50)	(3.04)	(0.74)	(2.99)	(0.46)	(2.95)	(.)	(.)
GAVI funding ($/capita)	0	0	0.095	0.00010	0.24	0.0026	0.40	0.057	0.59	0.074	0.65	0.077
	(0)	(0)	(0.15)	(0.0015)	(0.24)	(0.017)	(0.37)	(0.17)	(0.51)	(0.21)	(0.56)	(0.19)
GAVI funding ($/infant)	0	0	2.96	0.0038	7.80	0.24	12.3	3.35	19.3	3.33	21.1	3.54
	(0)	(0)	(4.16)	(0.054)	(7.22)	(1.56)	(11.2)	(9.22)	(16.7)	(8.70)	(17.5)	(7.63)
Probability of receiving GAVI (%)	0	0	75.4	0.60	83.8	3.53	88.5	19.5	90.4	19.8	88.7	24.7
	(0)	(0)	(43.2)	(7.72)	(36.9)	(18.5)	(32.3)	(39.9)	(29.8)	(40.1)	(32.0)	(43.4)
N	69	67	69	67	68	68	52	82	52	81	53	77

Note: Our sample is restricted to countries within 2 log points or less above or below the eligibility threshold. Phase 1 is defined as 2001-2005, phase 2: 2006-2010, phase 3: 2011, phase 4: 2012 and phase 5: 2013.

Table 3: Aid and income eligibility: first-stage regressions

	Bandwidth = 0.5		Bandwidth = 1		Bandwidth = 2	
	(1)	(2)	(3)	(4)	(5)	(6)
E	0.559***	0.591***	0.550***	0.503***	0.524***	0.520***
	(0.128)	(0.195)	(0.0867)	(0.131)	(0.0673)	(0.0897)
$\bar{Y}$	0.0916	0.143	-0.164*	-0.114	-0.117***	-0.223*
	(0.322)	(1.048)	(0.0901)	(0.353)	(0.0314)	(0.119)
$\bar{Y}^2$		-0.103		-0.0492		0.0557
		(2.008)		(0.328)		(0.0519)
$E \times \bar{Y}$	-0.242	0.0298	0.155	-0.208	-0.00737	0.200
	(0.495)	(1.474)	(0.130)	(0.550)	(0.0578)	(0.203)
$E \times \bar{Y}^2$		0.741		-0.236		0.000772
		(3.398)		(0.543)		(0.0937)
Constant	0.156*	0.151	0.206***	0.197**	0.180***	0.216***
	(0.0852)	(0.120)	(0.0635)	(0.0857)	(0.0450)	(0.0655)
Optimal b.w.						
Observations	435	435	1014	1014	1757	1757
R-squared	0.33	0.33	0.42	0.42	0.58	0.58

Note: The table reports estimates of the first-stage regression for the 2SLS fuzzy RD model. Each column presents a separate OLS regression, the dependent variable is the GAVI treatment dummy. Standard errors are clustered at the country level.

Table 4: Balance test: applying the RD to baseline data

	Vaccination Rate					Mortality Rate	
	(1) Hep B	(2) Hib	(3) DPT	(4) "Penta"	(5) MCV	(6) ≤ 5 yr	(7) ≤ 1 yr
Bandwidth = 2							
Linear	-16.2	5.6	-3.0	-4.5	-4.9	-8.5	-2.5
	(13.5)	(8.0)	(5.8)	(6.4)	(5.5)	(12.5)	(7.6)
Quadratic	-25.8	4.5	2.7	-6.2	0.2	-6.1	-4.6
	(19.0)	(11.9)	(8.2)	(8.9)	(7.6)	(16.6)	(10.8)
Obs.	136	136	136	136	136	136	136
No. of countries	136	136	136	136	136	136	136
Bandwidth = 1							
Linear	-18.9	5.7	1.3	-4.0	-2.0	-3.2	-3.7
	(17.7)	(10.8)	(7.6)	(8.5)	(7.2)	(16.0)	(10.1)
Quadratic	-33.3	-7.5	12.8	-9.4	13.9	-34.7	-18.7
	(24.5)	(14.5)	(10.5)	(10.9)	(9.9)	(20.9)	(13.9)
Obs.	81	81	81	81	81	81	81
No. of countries	81	81	81	81	81	81	81
Bandwidth = 0.5							
Linear	-43.0*	-9.8	12.5	-13.4	12.6	-31.9	-16.4
	(23.9)	(13.0)	(10.3)	(10.1)	(9.7)	(20.5)	(13.7)
Quadratic	-32.7	-33.4	16.0	-16.7	7.7	-11.0	-6.2
	(30.5)	(23.1)	(11.9)	(14.7)	(12.1)	(26.3)	(18.5)
Obs.	35	35	35	35	35	35	35
No. of countries	35	35	35	35	35	35	35

Note: The sample includes data points from 2000, before GAVI began disbursement in 2001. Each coefficient in the table (i.e., each row and each column) represents a separate regression. Each column uses a different dependent variable, listed in the column heading. Each row uses a different set of controls (i.e., polynomials of GNI and interactions with the GAVI treatment dummy). Standard errors are clustered at the country level.

Table 5: Impact estimates: intent-to-treat effect (ITT)

	Vaccination Rate							Mortality Rate		Spending ($)	
	(1) Hep B	(2) Hib	(3) DPT	(4) "Penta"	(5) Rota	(6) Pneumo	(7) MCV	(8) ≤ 5 yr	(9) ≤ 1 yr	(10) Total	(11) Gov't
Bandwidth = 2											
Linear	-8.4	4.3	-4.3	-2.8	-2.2	0.9	-5.3	9.7	8.0	-0.06	0.2
	(5.5)	(8.2)	(3.0)	(4.3)	(2.9)	(2.3)	(3.5)	(8.2)	(5.3)	(0.4)	(0.2)
Quadratic	-4.3	10.5	0.8	2.3	1.4	-2.6	0.3	14.7	9.8	0.5	-0.1
	(6.4)	(12.2)	(3.9)	(6.0)	(3.0)	(2.9)	(4.7)	(10.9)	(7.2)	(0.5)	(0.2)
Obs.	1754	1754	1754	1754	1754	1754	1754	404	404	795	823
No. of countries	149	149	149	149	149	149	149	147	147	136	138
Bandwidth = 1											
Linear	-0.4	10.2	1.4	3.7	0.9	1.1	-0.05	8.5	6.4	1.2	-0.04
	(5.6)	(11.4)	(3.6)	(5.4)	(2.7)	(2.4)	(4.5)	(9.6)	(6.2)	(1.0)	(0.2)
Quadratic	-7.7	0.7	3.4	-1.2	5.5	-2.3	6.2	-9.3	-5.6	-1.4	0.2
	(8.8)	(16.8)	(5.0)	(8.0)	(3.8)	(3.2)	(5.3)	(12.1)	(8.6)	(1.3)	(0.2)
Obs.	1011	1011	1011	1011	1011	1011	1011	228	228	455	476
No. of countries	101	101	101	101	101	101	101	98	98	87	90
Bandwidth = 0.5											
Linear	-5.3	-2.9	4.9	-1.1	-0.4	-2.4	7.1	-16.0	-9.2	0.05	0.02
	(9.0)	(16.7)	(5.5)	(8.2)	(3.5)	(3.0)	(5.7)	(10.8)	(8.0)	(0.2)	(0.2)
Quadratic	5.5	-28.5	11.3	-3.9	0.8	-5.2	8.4	-5.2	-3.2	0.1	0.03
	(13.4)	(22.7)	(7.1)	(11.3)	(9.9)	(4.6)	(6.9)	(19.2)	(13.3)	(0.5)	(0.4)
Obs.	435	435	435	435	435	435	435	101	101	211	213
No. of countries	64	64	64	64	64	64	64	63	63	50	50

Note: Each coefficient in the table (i.e., each row and each column) represents the ITT treatment effect from a separate regression. Each column uses a different dependent variable, listed in the column heading. Each row uses a different set of controls (i.e., polynomials of GNI and interactions with the GAVI treatment dummy). Standard errors are clustered at the country level.

Table 6: Impact estimates: intent-to-treat effect (ITT) - controlling for baseline vaccination rate

	Vaccination Rate							Mortality Rate		Spending ($)	
	(1) Hep B	(2) Hib	(3) DPT	(4) "Penta"	(5) Rota	(6) Pneumo	(7) MCV	(8) $\leq$ 5 yr	(9) $\leq$ 1 yr	(10) Total	(11) Gov't
Bandwidth = 2											
Linear	-0.5	1.7	-0.1	2.1	-2.2	0.9	-0.8	3.5	2.3	-0.06	0.2
	(5.2)	(7.9)	(2.1)	(4.2)	(2.9)	(2.3)	(2.6)	(3.2)	(1.8)	(0.4)	(0.2)
Quadratic	5.9	9.1	4.0	8.5	1.6	-2.3	4.1	4.8	3.1	0.4	-0.1
	(6.5)	(11.9)	(3.0)	(5.9)	(3.0)	(2.9)	(3.6)	(3.7)	(2.0)	(0.5)	(0.2)
Obs.	1737	1737	1737	1737	1737	1737	1737	399	399	790	818
No. of countries	144	144	144	144	144	144	144	143	143	134	136
Bandwidth = 1											
Linear	7.2	4.5	4.0	7.7	1.0	1.2	3.5	-0.3	0.8	1.2	-0.04
	(5.9)	(10.8)	(2.6)	(5.4)	(2.7)	(2.4)	(3.4)	(2.7)	(1.5)	(1.0)	(0.2)
Quadratic	2.0	-0.7	2.5	3.6	5.6	-2.2	5.4	3.3	1.7	-1.4	0.2
	(8.1)	(16.0)	(4.0)	(7.4)	(3.8)	(3.2)	(4.5)	(3.3)	(2.1)	(1.3)	(0.2)
Obs.	1002	1002	1002	1002	1002	1002	1002	226	226	451	472
No. of countries	99	99	99	99	99	99	99	97	97	86	89
Bandwidth = 0.5											
Linear	3.7	-0.5	3.7	4.1	-0.1	-2.1	5.9	0.7	0.4	0.07	0.03
	(9.3)	(15.6)	(4.1)	(8.1)	(3.5)	(3.0)	(4.4)	(3.2)	(2.2)	(0.2)	(0.2)
Quadratic	9.7	-20.5	3.5	-3.2	1.0	-5.0	3.3	6.7	3.2	0.08	0.01
	(12.7)	(20.7)	(5.7)	(9.9)	(9.9)	(4.5)	(6.3)	(4.1)	(2.5)	(0.5)	(0.4)
Obs.	432	432	432	432	432	432	432	100	100	209	211
No. of countries	62	62	62	62	62	62	62	62	62	49	49

Note: Each coefficient in the table (i.e., each row and each column) represents the ITT treatment effect from a separate regression, with baseline vaccination rate added as a control. Each column uses a different dependent variable, listed in the column heading. Each row uses a different set of controls (i.e., polynomials of GNI and interactions with the GAVI treatment dummy). Standard errors are clustered at the country level.

Table 7: Impact estimates: effect of treatment on the treated (TOT)

	Vaccination Rate							Mortality Rate		Spending ($)	
	(1) Hep B	(2) Hib	(3) DPT	(4) "Penta"	(5) Rota	(6) Pneumo	(7) MCV	(8) $\leq$ 5 yr	(9) $\leq$ 1 yr	(10) Total	(11) Gov't
Bandwidth = 2											
Linear	-17.7	10.1	-8.8	-5.5	-4.6	2.1	-10.9	19.6	16.3	-0.1	0.3
	(11.4)	(16.7)	(6.0)	(8.6)	(5.9)	(4.6)	(7.2)	(16.6)	(10.5)	(0.7)	(0.3)
Quadratic	-7.0	19.9	2.3	5.1	3.4	-4.6	1.6	26.4	18.8	0.3	0.02
	(13.5)	(23.8)	(8.0)	(11.8)	(6.0)	(5.7)	(9.4)	(24.4)	(16.0)	(0.4)	(0.3)
Obs.	1754	1754	1754	1754	1754	1754	1754	404	404	795	823
No. of countries	149	149	149	149	149	149	149	147	147	136	138
Bandwidth = 1											
Linear	1.8	17.8	3.2	7.6	2.1	1.3	0.07	12.4	11.0	1.2	-0.08
	(11.4)	(22.0)	(7.3)	(10.6)	(5.4)	(4.6)	(8.9)	(22.3)	(14.2)	(1.1)	(0.2)
Quadratic	-19.5	-3.0	8.1	-4.8	13.5	-6.3	15.9	-33.7	-21.0	-4.5	0.6
	(26.4)	(40.0)	(11.9)	(20.4)	(10.5)	(9.0)	(13.2)	(38.8)	(27.0)	(7.0)	(1.0)
Obs.	1011	1011	1011	1011	1011	1011	1011	228	228	455	476
No. of countries	101	101	101	101	101	101	101	98	98	87	90
Bandwidth = 0.5											
Linear	-12.8	-7.9	10.1	-3.5	-1.8	-5.3	14.8	-33.5	-19.3	0.07	0.01
	(21.0)	(37.7)	(11.7)	(18.6)	(8.4)	(7.0)	(12.2)	(32.2)	(21.2)	(0.3)	(0.3)
Quadratic	17.5	-66.1	23.4	-8.4	1.2	-7.9	15.6	2.0	0.9	-0.009	-0.2
	(31.4)	(75.2)	(15.6)	(27.4)	(17.7)	(11.2)	(14.2)	(50.8)	(32.8)	(0.8)	(1.1)
Obs.	435	435	435	435	435	435	435	101	101	211	213
No. of countries	64	64	64	64	64	64	64	63	63	50	50

Note: Each coefficient in the table (i.e., each row and each column) represents the "treatment on the treated" (TOT) effect from a separate 2SLS regression, aka fuzzy RD model. Each column uses a different dependent variable, listed in the column heading. Each row uses a different set of controls (i.e., polynomials of GNI and interactions with the GAVI treatment dummy). The GAVI treatment dummy and its interactions with log baseline GNI are instrumented with the GAVI *eligibility* dummy and its equivalent interactions. Standard errors are clustered at the country level.

Table 8: Testing for fungibility and waste

	(1) Hep B	(2) Hib	(3) DPT	(4) "Penta"	(5) Rota	(6) Pneumo	(7) MCV	(8) Spending
Bandwidth = 2								
Linear	-95.5***	-7.1	-36.4***	-46.3***	-2.1	2.9	-10.9	-0.7
	(15.5)	(15.9)	(13.3)	(10.7)	(5.5)	(4.0)	(7.2)	(0.7)
Quadratic	-61.2**	3.9	-15.6	-24.3	4.2	-2.6	2.1	-0.4
	(24.4)	(21.8)	(17.8)	(15.1)	(5.9)	(4.0)	(9.3)	(0.3)
Obs.	1754	1754	1754	1754	1754	1754	1742	795
No. of countries	149	149	149	149	149	149	149	136
Bandwidth = 1								
Linear	-56.1***	-6.2	-29.8*	-30.7**	3.8	2.2	0.7	0.4
	(21.3)	(19.8)	(16.7)	(13.1)	(5.3)	(3.0)	(8.8)	(1.1)
Quadratic	-84.6*	14.8	44.4	-8.5	11.6	-1.5	15.9	-5.1
	(46.2)	(39.9)	(33.9)	(28.0)	(10.5)	(4.6)	(13.2)	(6.9)
Obs.	1011	1011	1011	1011	1011	1011	1004	455
No. of countries	101	101	101	101	101	101	101	87
Bandwidth = 0.5								
Linear	-72.6*	-15.7	32.7	-18.5	4.4	-1.3	14.6	-0.5**
	(40.0)	(35.8)	(25.6)	(25.2)	(5.6)	(1.5)	(12.3)	(0.2)
Quadratic	-53.8	-68.4	22.2	-33.3	6.9	-2.4	15.5	-0.7
	(57.3)	(82.7)	(26.4)	(39.3)	(9.7)	(3.0)	(14.4)	(0.7)
Obs.	435	435	435	435	435	435	434	211
No. of countries	64	64	64	64	64	64	64	50

Note: This table tests the claim that GAVI-delivered dose series were additional and did not offset the number of doses supplied by other sources. Coefficients less than zero imply fungibility and/or waste. Each coefficient in the table (i.e., each row and each column) represents the "treatment on the treated" effect from a separate 2SLS regression. Each column uses a different dependent variable, listed in the column heading. Each row uses a different set of controls (i.e., polynomials of GNI and interactions with the GAVI treatment dummy). The GAVI treatment dummy and its interactions with log baseline GNI are instrumented with the GAVI *eligibility* dummy and its equivalent interactions. Standard errors are clustered at the country level.

Table 9: Testing for waste

	(1) Hep B	(2) Hib	(3) DPT	(4) "Penta"	(5) Rota	(6) Pneumo
Bandwidth = 2						
Linear	-16.4***	-4.2**	0.003	-0.8	-1.3	-0.9
	(5.5)	(2.0)	(0.8)	(1.1)	(1.1)	(0.5)
Quadratic	-4.0	-5.7*	-0.09	-0.03	-2.7	-1.3
	(9.0)	(3.0)	(0.5)	(0.8)	(2.2)	(1.1)
Obs.	1754	1754	1754	1754	1754	1754
No. of countries	149	149	149	149	149	149
Bandwidth = 1						
Linear	-3.1	-6.4**	-0.3	-0.3	-2.4	-1.5
	(8.9)	(2.7)	(0.3)	(0.6)	(2.0)	(1.0)
Quadratic	-19.9	-5.0	1.7	1.9	-4.0	-2.0
	(14.8)	(4.6)	(1.8)	(1.7)	(3.4)	(1.6)
Obs.	1011	1011	1011	1011	1011	1011
No. of countries	101	101	101	101	101	101
Bandwidth = 0.5						
Linear	-14.3	-6.4	1.7	1.7	-3.6	-1.8
	(11.9)	(4.2)	(1.7)	(1.4)	(3.1)	(1.4)
Quadratic	-18.5	-7.5	-0.5	-0.6	-2.3	-0.9
	(13.8)	(5.3)	(1.0)	(1.1)	(2.1)	(1.2)
Obs.	435	435	435	435	435	435
No. of countries	64	64	64	64	64	64

Note: This table tests the null hypothesis that all GAVI aid translated directly into a marginal increase in vaccination rates. Coefficients less than zero imply waste. Each coefficient in the table (i.e., each row and each column) represents the "treatment on the treated" effect from a separate regression. Each column uses a different dependent variable, listed in the column heading. Each row uses a different set of controls (i.e., polynomials of GNI and interactions with the GAVI treatment dummy). Standard errors are clustered at the country level.

Table 10: Decomposition: Where did free vaccines go?

	HepB	Hib	DPT	Penta	Rota	Pneumo	Measles
(A) GAVI deliveries	58.4	29.6	35.5	29.2	4.2	3.7	-0.5
	(11.61)	(6.60)	(8.39)	(6.56)	(1.62)	(1.90)	(1.15)
≈ (B) GAVI impact	1.8	17.8	3.2	7.6	2.1	1.3	0.07
	(11.4)	(22.0)	(7.3)	(5.4)	(5.4)	(4.6)	(8.9)
+ (C) Fungibility	53.0	0	29.5	30.4	0	0	0
	(21.3)	(.)	(16.7)	(.)	(.)	(.)	(.)
+ (D) Waste	3.1	6.4	0.3	0.3	2.4	1.5	0
	(8.9)	(2.7)	(0.3)	(2.0)	(2.0)	(1.0)	(.)

Note: This table presents estimates of vaccine deliveries, impact, waste and fungibility at the eligibility threshold. GAVI deliveries are measured as the percent of the population under 1 provided with a full series of doses by GAVI. Estimates of impact attributable to GAVI, fungibility and waste are drawn from the linear models, bandwidth 1, in Table 7, Table 8 and Table 9, respectively. Estimates of fungibility and waste are capped at 0.

Figure 1: Immunization rates, 2000 and 2013

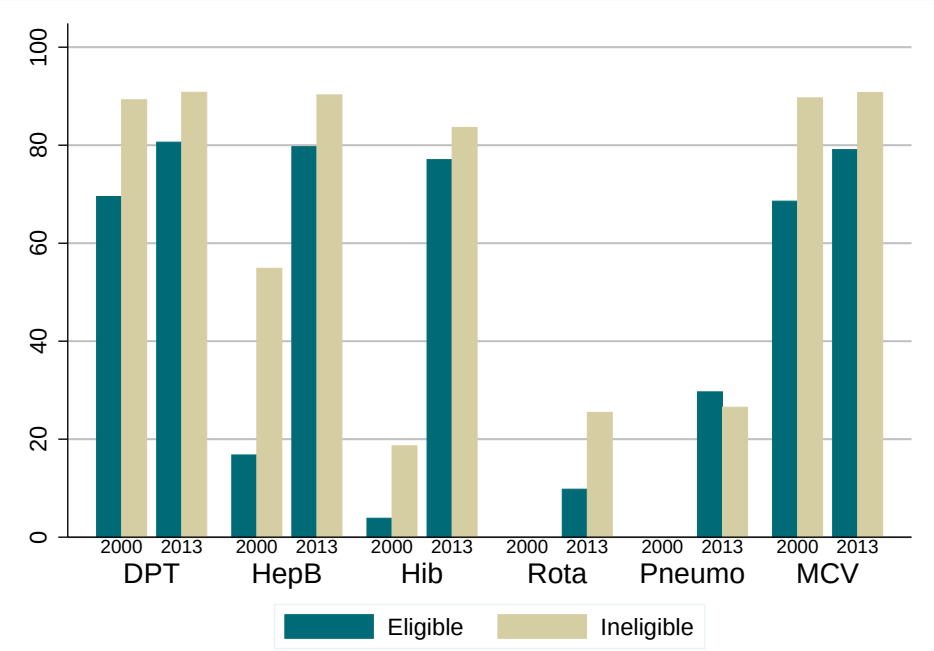

Figure 2: Government and non-government vaccine financing, 2000 and 2012

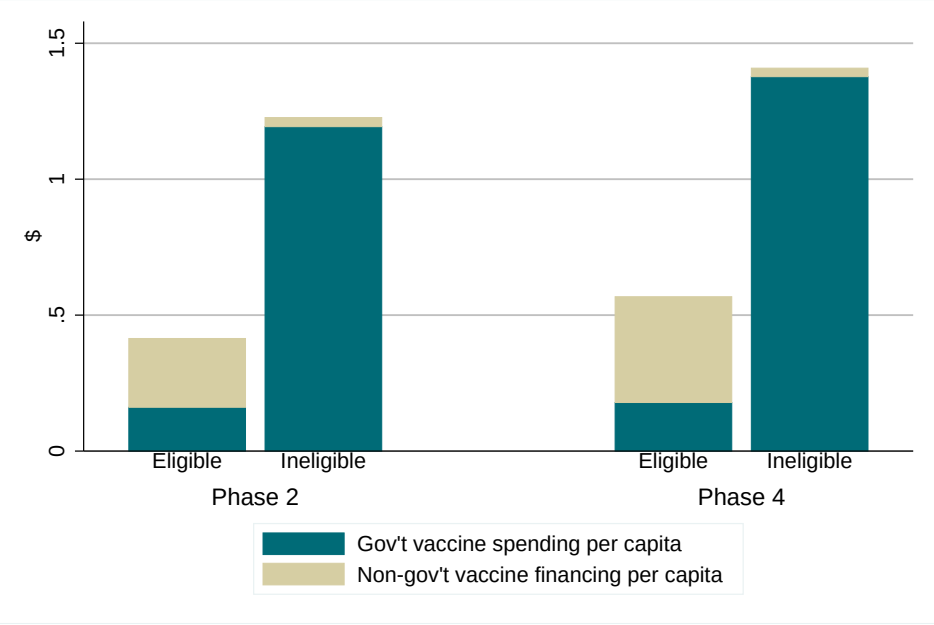

Figure 3: GAVI's adherence to its income eligibility threshold

(a) Funding amount per infant

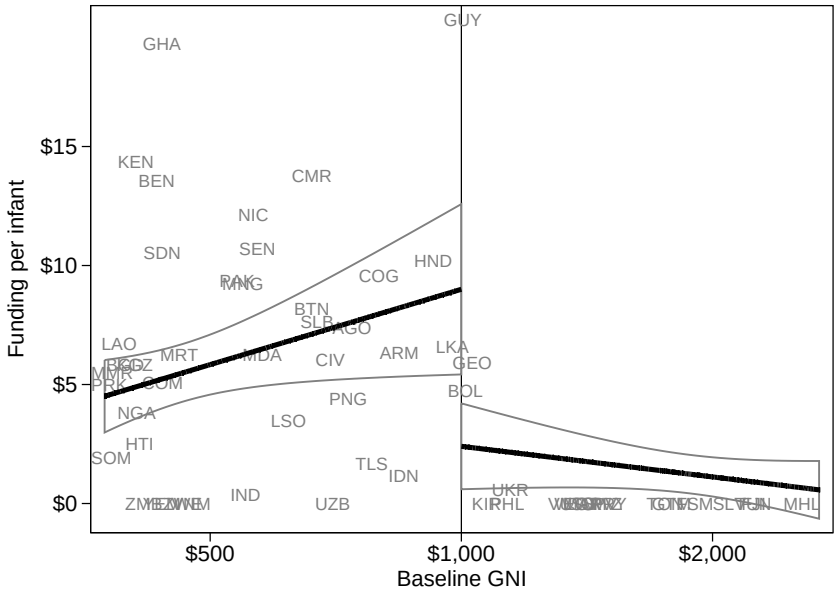

(b) Probability of receiving funding

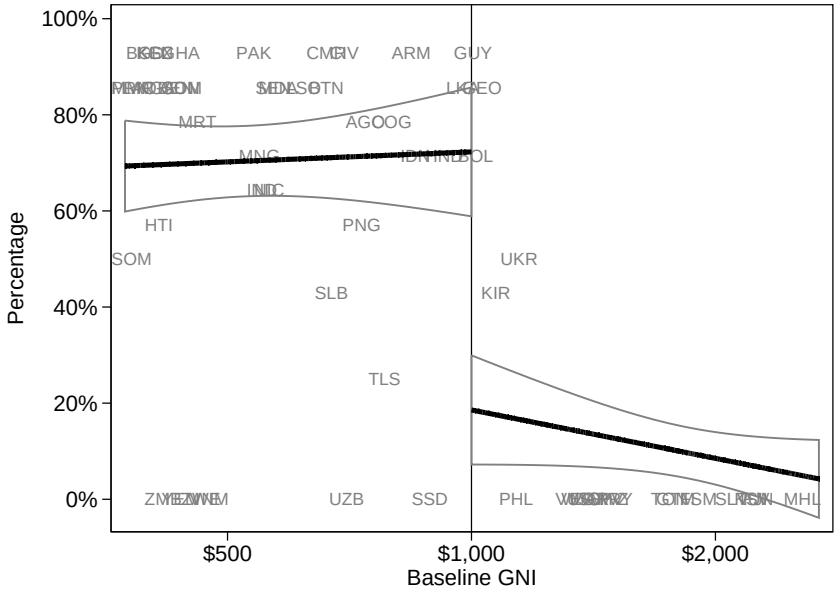

The figures measure GAVI's adherence to its own eligibility rules. Each graph presents a local polynomial regression of GAVI funding (dollars per child in top graph, and a binary indicator for receipt of funding in bottom graph) on countries' 1998 log per capita GNI, run separately on each side of the administrative cut-off at $1,000. Gray lines denote the 95% confidence interval around the regression line. Scatterplots show average values in discrete bins defined by GNI, each bin spanning 0.2 log points.

Figure 4: Vaccination rates, 2000-2013, by GNI

(a) Hepatitis B

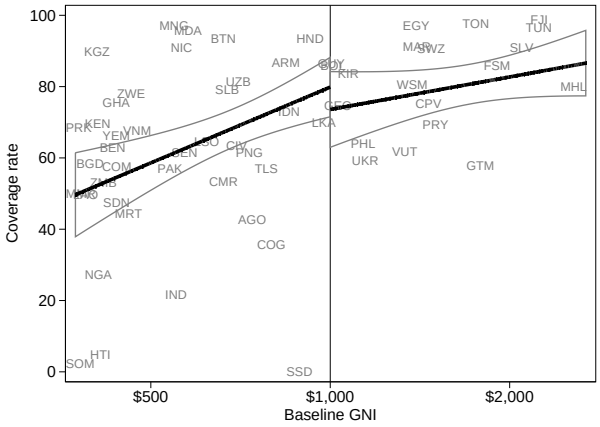

(b) DPT

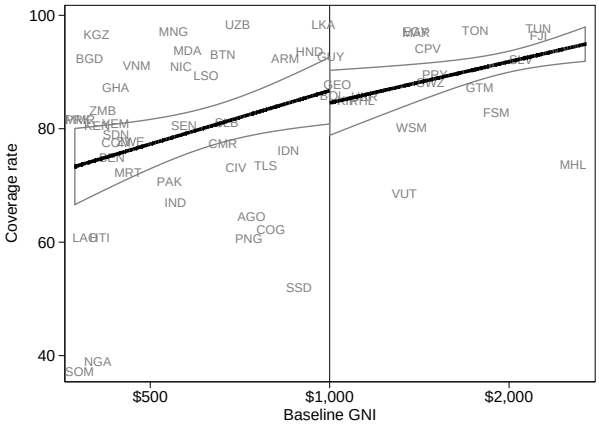

(c) Hib

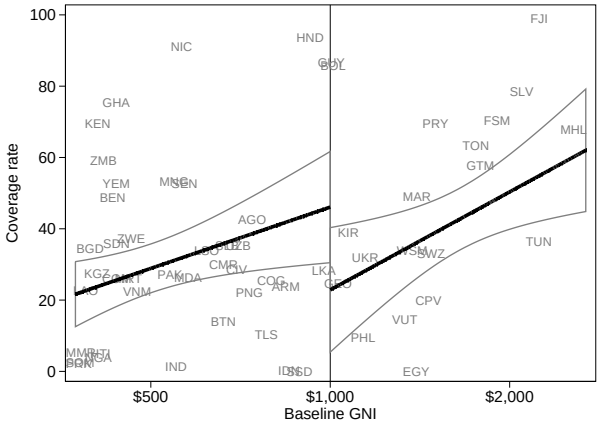

(d) Rota

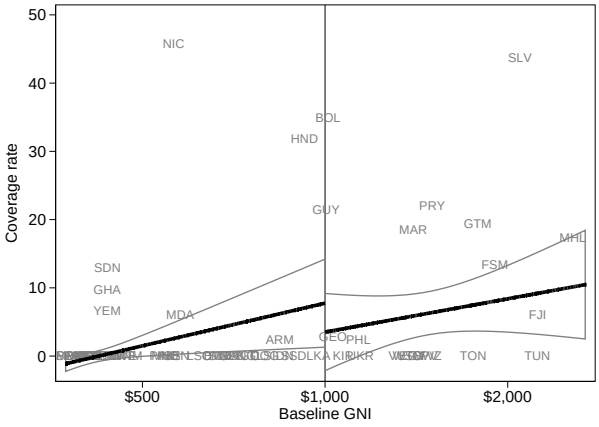

(e) Pneumococcal

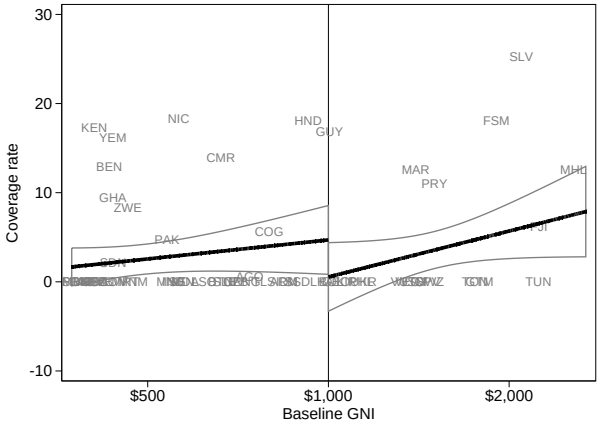

(f) Measles (MCV)

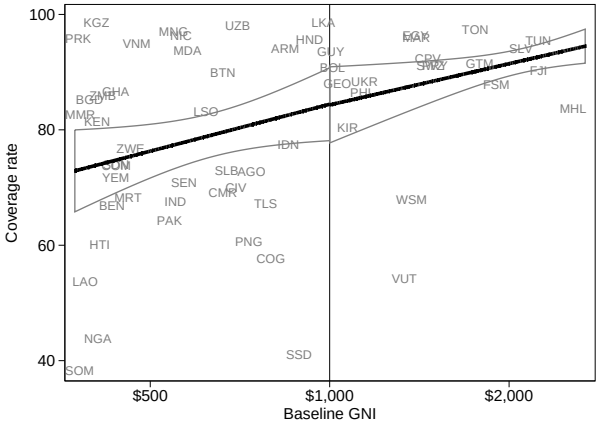

Figures show a linear prediction plot of vaccine coverage on countries' baseline log per capita GNI, run separately on each side of the administrative cut-off at $1,000. The solid black lines use the actual data series. Gray lines denote the 95% confidence interval around the regression line.

Figure 5: GAVI-purchased dose series and vaccination rates, 2001-2013

(a) Hepatitis B

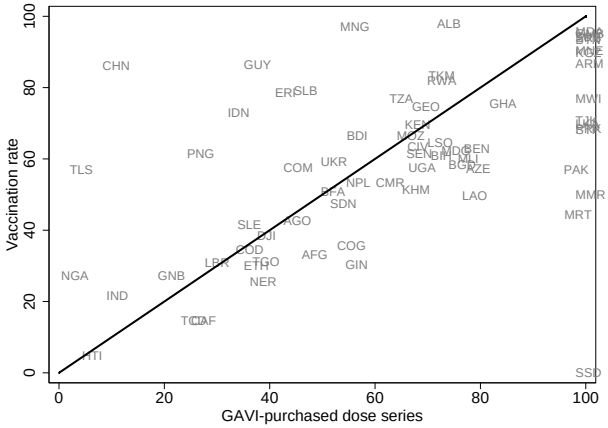

(b) DPT

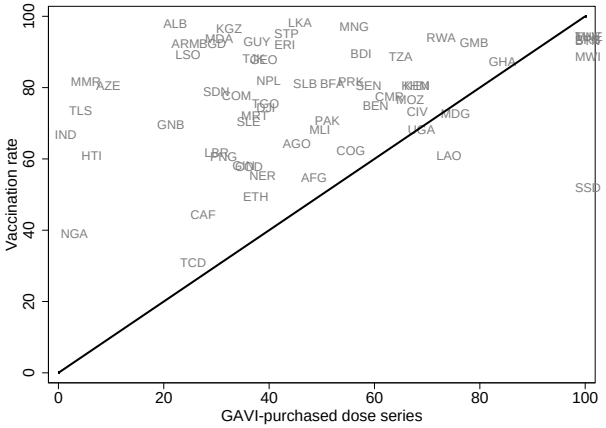

(c) Hib

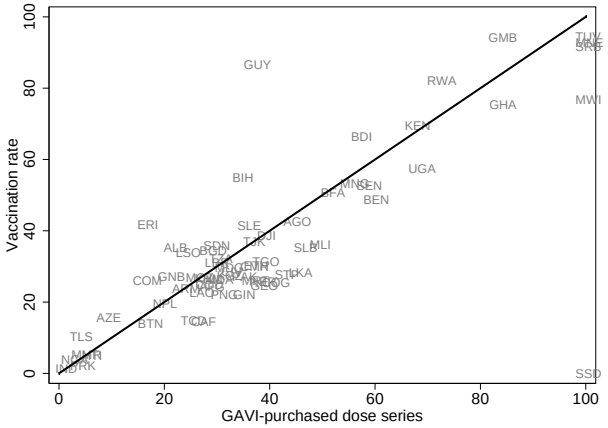

(d) "Penta"

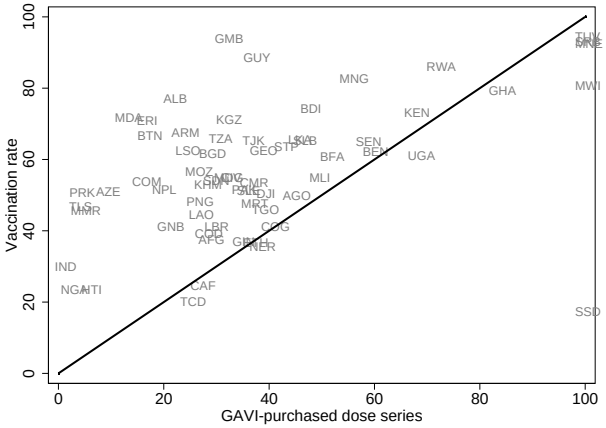

(e) Rota

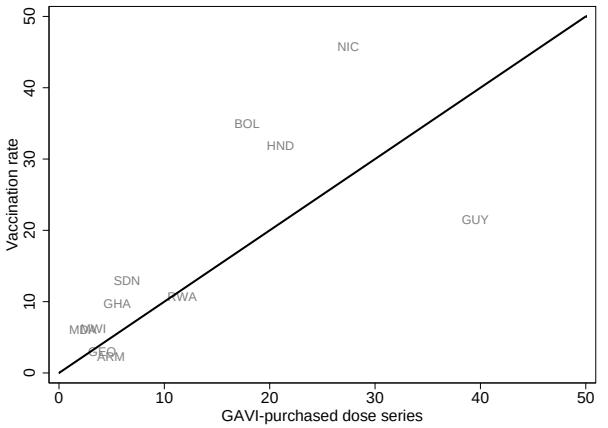

(f) Pneumococcal

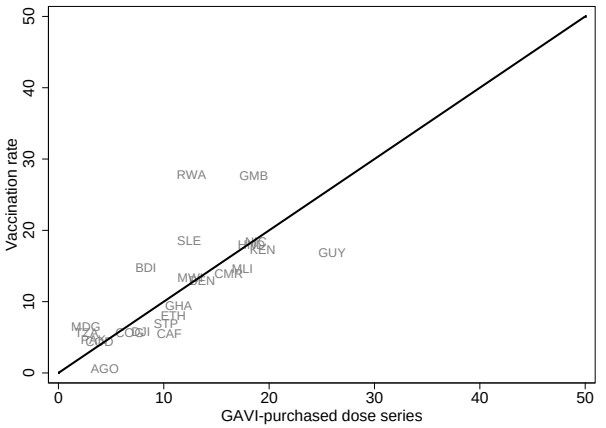

Scatterplots show the average number of complete dose series purchased by GAVI, as a percent of the population under 1, and the average vaccination rate by country during the period 2001-2013 for countries receiving vaccines for at least one of those years. The solid black line is a 45-degree line. Points below the line represent countries which received more doses from GAVI than the number of children under 1 vaccinated on average. Percent of dose series purchased by GAVI is capped at 100%.

Figure 6: Testing levels of fungibility, 2000-2013

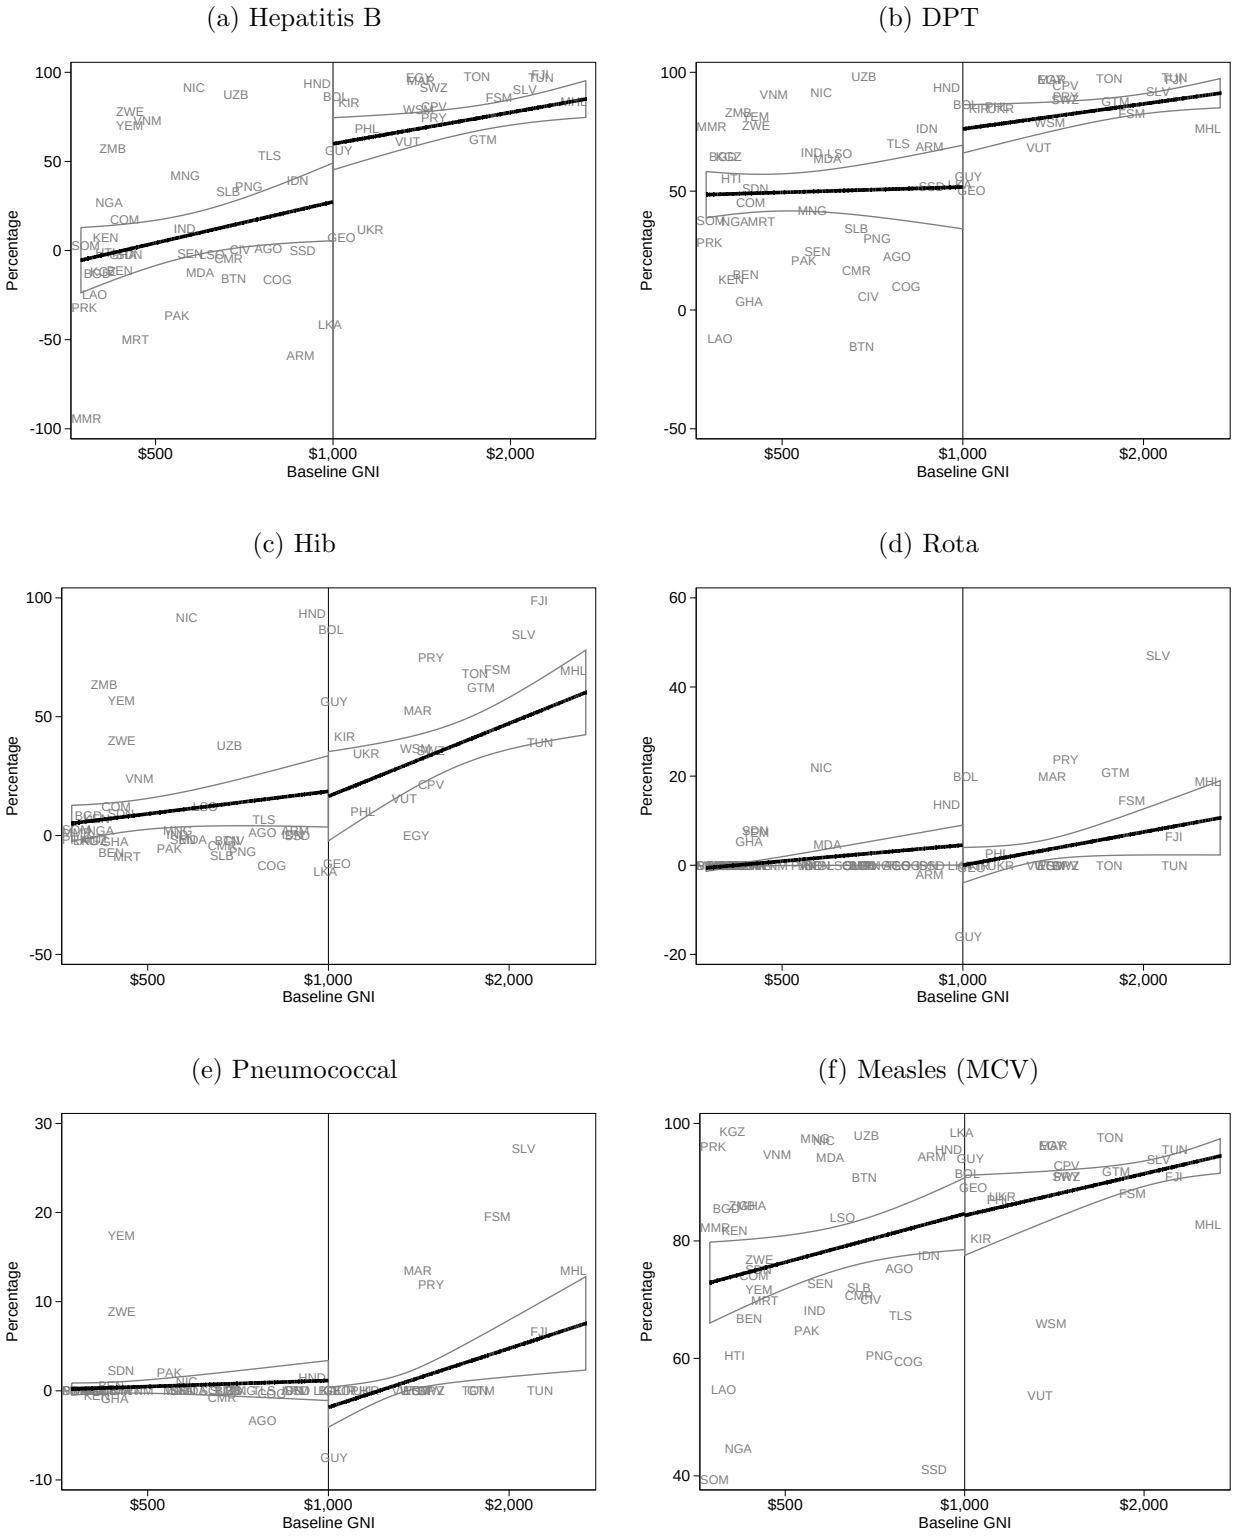

Figures show a linear prediction plot of vaccine coverage on countries' baseline log per capita GNI, run separately on each side of the administrative cut-off at $1,000. The solid black lines subtract GAVI doses from actual vaccination series: i.e., the counterfactual to GAVI if all GAVI aid represents a marginal increase in vaccination rates. Gray lines denote the 95% confidence interval around the regression line.

Figure 7: Child mortality and vaccination funding

(a) Mortality 0-5 years old, 2001-12

(b) Mortality 0-1 year old, 2001-12

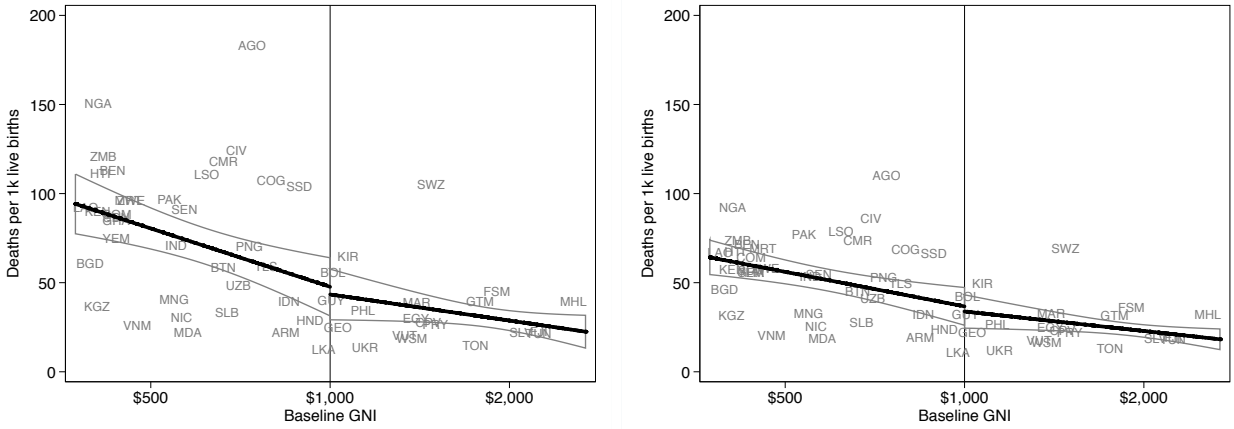

(c) Governments' own expenditure

(d) Total expenditure (including foreign aid)

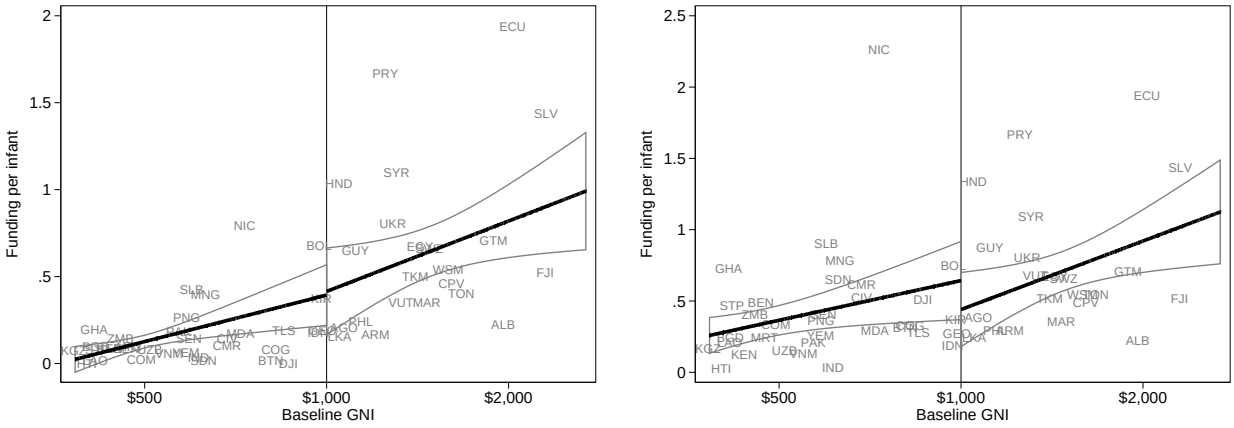

Panel (a) shows a linear prediction plot of under-five mortality (deaths per 1,000 live births) on countries' 1998 log per capita GNI, run separately on each side of the administrative cut-off at $1,000. Gray lines denote the 95% confidence interval around the regression line. Scatterplots show average values in discrete bins defined by GNI, each bin spanning 0.2 log points. Panel (b) repeats this exercise using annual changes in the mortality rate on the y-axis. Panel (c) shows results for the log of per capita expenditure on vaccines by national governments. Fungibility would imply a negative GAVI treatment effect. Panel (d) shows results for the log of total per capital expenditure on vaccines, including domestic spending and foreign aid. Anything short of a positive treatment effect equal to the amount of GAVI aid could be interpreted as evidence of fungibility.